

ABSTRACT

GADDY, MELISSA ROSE. Optimization of Radiotherapy Treatments with Spatiotemporal Fractionation Schemes. (Under the direction of Dr. Dávid Papp).

In this thesis, we develop mathematical optimization techniques to improve external beam photon radiotherapy treatments. We begin by revisiting a treatment planning optimization method for arc therapy planning called VMERGE and propose two modifications to the algorithm. First, we incorporate an explicit constraint on the treatment time during the treatment planning optimization and implement an efficient algorithm that exploits the structure of the problem to optimize a solution. Next, we substitute a heuristic in the original VMERGE algorithm with a greedy strategy to obtain a treatment that improves the trade-off of treatment efficacy versus treatment time. We demonstrate the benefit of these two modifications on a clinical prostate and paraspinal case.

Next, we optimize spatiotemporally fractionated treatments by incorporating a mathematical model of a tissue's response to radiation over a series of repeated treatments into the treatment planning optimization. In spatiotemporal plans, the radiation dose distribution is allowed to be different on each treatment day, and these plans achieve a balance of the potentially competing objectives of delivering high doses to the target volumes on a single treatment day while maintaining consistent low doses to the healthy tissues. Spatiotemporal plans exhibit an improvement in healthy organ sparing over treatments optimized with a conventional treatment schedule.

The optimization models for spatiotemporal fractionation are nonconvex, so the spatiotemporally fractionated treatment plans are only certified to be locally optimal. We formulate a semidefinite programming relaxation of the nonconvex optimization problem to derive a bound on the maximum benefit of spatiotemporal fractionation over conventional treatment schedules. We show that the locally optimal spatiotemporally fractionated treatment plans attain a large percentage of the maximum possible improvement in healthy tissue sparing.

The degradation of plan quality in the presence of patient setup uncertainty is a more pressing concern for spatiotemporal fractionation than for conventional treatments, and the conventional

methods to mitigate setup uncertainty cannot be used for spatiotemporal planning. In this work, we utilize stochastic optimization techniques to incorporate random patient setup uncertainty into the optimization problems for spatiotemporally fractionated plans and demonstrate that even in the presence of patient setup uncertainty, spatiotemporal plans exhibit an improvement in normal tissue sparing over uniformly fractionated plans while maintaining robust target coverage.

Optimization of Radiotherapy Treatments with Spatiotemporal Fractionation Schemes

by
Melissa Rose Gaddy

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Applied Mathematics

Raleigh, North Carolina

2019

APPROVED BY:

Mansoor A. Haider

Arvind K. Saibaba

Julie S. Ivy

Dávid Papp
Chair of Advisory Committee

BIOGRAPHY

Melissa Gaddy earned a Bachelor of Arts in applied mathematics and a Bachelor of Science in computer science from Wofford College in 2014. She earned a Master of Science in applied mathematics from North Carolina State University in 2016.

ACKNOWLEDGEMENTS

I would like to express my gratitude to my advisor Dávid Papp for his guidance throughout this research. I greatly appreciate his patience, generous giving of his time, and careful critiques of my written work. I would also like to thank Sercan Yıldız for his collaboration and Jan Unkelbach for his clinical expertise. Finally, a loving thank-you to my family, for always saying you are proud of me.

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
Chapter 1 Introduction	1
1.1 External beam radiotherapy	1
1.2 Radiotherapy treatment modalities	3
1.2.1 Fixed beam intensity-modulated radiotherapy (IMRT)	3
1.2.2 Photon arc therapy (VMAT)	4
1.2.3 Intensity-modulated proton therapy (IMPT)	4
1.3 Optimization of radiotherapy treatments	6
1.3.1 Clinical objectives	8
1.3.2 A mathematical model of fluence map optimization	10
1.4 Fractionation of radiotherapy treatments	15
1.4.1 Spatiotemporal fractionation	15
1.5 Uncertainty in radiotherapy treatments	16
1.5.1 Sources of uncertainty in radiotherapy	16
1.5.2 Optimization under uncertainty	17
1.6 The contribution of this thesis	17
Chapter 2 Improving photon arc therapy delivery	19
2.1 Introduction	20
2.2 Two improvements to the VMERGE algorithm	22
2.2.1 A greedy approach to merging	22
2.2.2 Regularization and treatment time constraints	23
2.3 Optimizing a treatment plan with a constraint on the treatment time	25
2.3.1 Decomposition of the inner problem	28
2.3.2 The subproblem for each beamlet row	30
2.3.3 Linear-time algorithm to solve the reduced subproblem	32
2.4 Numerical experiments with clinical cases	33
2.4.1 Experimental setup	33
2.4.2 Results	35
2.5 Discussion	41
Chapter 3 Optimizing spatiotemporal plans	42
3.1 Introduction	42
3.1.1 The linear-quadratic model for biologically effective dose	42
3.1.2 Relation to prior works	44
3.2 A mathematical model for optimal spatiotemporal fractionation	46
3.2.1 Nonuniform fractionation using the BED model	46
3.2.2 Uniform fractionation using the BED model	48
3.2.3 The constrained nonuniform model	50

3.3	Application to liver tumors	51
3.3.1	Experimental setup	51
3.3.2	Results	54
3.4	Discussion	57
Chapter 4	Bounding the achievable benefit of spatiotemporal fractionation	62
4.1	Semidefinite programming relaxations of general QCQPs	63
4.2	A mathematical model for bounding benefit	65
4.2.1	Bounding the maximum benefit in a given objective	69
4.2.2	Solution methods	70
4.3	Results for clinical liver cases	71
4.4	Discussion	73
Chapter 5	Robustness in spatiotemporal planning	74
5.1	Introduction	75
5.2	Treatment plan optimization for spatiotemporal fractionation under uncertainty . . .	76
5.2.1	Modeling dose delivery uncertainty	76
5.2.2	Stochastic optimization of uniformly fractionated treatments	77
5.2.3	Stochastic optimization of spatiotemporal treatments	79
5.3	Modeling random patient setup uncertainty	80
5.3.1	Probability distributions of patient position	80
5.3.2	Computing the expected penalty values	81
5.4	Numerical experiments with clinical liver cases	84
5.4.1	Clinical liver cases and prescriptions	84
5.4.2	Probability distributions of patient shifts	85
5.4.3	Optimizing robust uniformly fractionated reference plans	86
5.4.4	Optimizing robust spatiotemporal plans	87
5.4.5	Computing equieffective dose	90
5.5	Results for clinical liver cases	90
5.6	Discussion	97
Chapter 6	Conclusions and open problems	99
References	101
Appendices	113
Appendix A	Supplementary figures for Chapter 3	114
Appendix B	Supplementary figures for Chapter 5	119

LIST OF TABLES

Table 3.1	Summary of mean liver BED reductions from spatiotemporal fractionation and sparing factors for each of the five cases. The sparing factor is a value that determines the dependence of the optimal fractionation schedule of a fixed dose distribution on the patient geometry. The remarkable benefit seen in Case 2 agrees with the fact that the sparing factor is substantially lower in this case than in the other cases.	55
Table 4.1	Summary of mean liver BED reductions from spatiotemporal fractionation and lower bounds for mean liver BED. The “gap closed” values provide a measure of how close the local optimal solutions are to achieving the lower bound on the mean liver BED; see Eq. (4.6) for the definition.	72
Table 5.1	Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 2-voxel shifts in either axis of the transverse plane. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to 2, where there is an equal probability of shifting to the 25 points around the nominal position. The voxel sizes are 2.54 mm \times 2.54 mm for Cases 1-4 and 2.18 mm \times 2.18 mm for Case 5.	91
Table 5.2	Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 1-voxel shifts in either axis of the transverse plane. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to $2/3$, where there is an equal probability of shifting to the 9 points around the nominal position. The voxel sizes are 2.54 mm \times 2.54 mm for Cases 1-4 and 2.18 mm \times 2.18 mm for Case 5.	91
Table 5.3	Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 1-voxel shifts in either axis of the transverse plane with the full-resolution dose grid. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to $2/3$, where there is an equal probability of shifting to the 9 points around the nominal position. The voxel sizes are 1.27 mm \times 1.27 mm for Cases 1-4 and 1.09 mm \times 1.09 mm for Case 5.	92

LIST OF FIGURES

Figure 1.1	(a) The equipment setup for radiation treatment. The patient lies on the flat table, and the linear accelerator rotates around the patient on a gantry. (b) Close-up image of a multi-leaf collimator, which is positioned in front of the source of the radiation beam in the linear accelerator to control the beam's size and shape. Images courtesy of Varian Medical Systems, Inc., Palo Alto, CA. Copyright 2019. All rights reserved.	2
Figure 1.2	This graphic represents an example of a fixed-beam intensity-modulated radiotherapy (IMRT) treatment [75]. Three photon beams with different intensity levels are delivered through the three apertures on the left in succession from the same beam angle, and the resulting beam on the right has an inhomogeneous intensity that is the sum of the intensities in each of the beams on the left. Reprinted by permission from the publisher.	5
Figure 1.3	(a) A comparison of the depth-dose curves for photon and proton beams [60]. Reprinted with permission. © 2014 American Society of Clinical Oncology. All rights reserved. (b) A comparison of the dose absorption within a patient using photon and proton beams [86]. (The acronym IMXT is intensity-modulated x-ray therapy, which is the same as IMRT.) Consistent with the depth-dose curves in (a), the dose distribution from the photon treatment delivers dose in front of and behind the tumor, while the proton beam avoids healthy tissues behind the tumor. Reprinted with permission from the publisher.	7
Figure 1.4	A screenshot of the open-source treatment planning system matRad [104]. Within the interface of a treatment planning system, a dosimetrist can complete all the steps of the treatment planning process, which includes specifying prescription values for each structure, calling the optimization solver, and evaluating treatment plan quality.	9
Figure 1.5	An example of a dose volume histogram (DVH) used to evaluate the quality of a treatment plan for a prostate case. Each curve represents a different structure within the patient, and for a fixed amount of dose, the DVH displays the fraction of the structure which receives at least that amount of dose. The prostate and seminal vesicles are the target volumes with a prescription dose of 70 Gy, and the remaining structures are OARs.	11
Figure 1.6	A pictorial representation of the setup of an IMRT problem [75]. The oval figure at the bottom represents the patient, which is divided into a grid of voxels. In the gantry head, the beam aperture is divided into a grid of beamlets, and the beamlet intensities are the decision variables in the fluence map optimization problem. Reprinted by permission from the publisher.	12

Figure 2.1	Representative slices of CT images for the clinical cases in this chapter. (a) The prostate case from the CORT dataset. The visible contoured OARs in this slice are the bladder (red), left and right femoral heads (light blue and brown), and the rectum (orange). The two PTVs are contoured in blue and yellow. (b) The paraspinal case. The PTV (contoured in yellow) is wrapped around the spinal cord, which is contoured in red.	34
Figure 2.2	Comparison of the original VMERGE method and the greedy variant on the prostate case, without SPG-regularization. Panel (a): objective function value versus treatment time. Panel (b): standard error of PTV dose (Eq. (2.17)) versus treatment time for the two PTVs. Panel (c): mean dose to the OARs with the original VMERGE method. Panel (d): mean dose to the OARs with the greedy variant. The data points with the highest treatment time represent the 180-beam IMRT plan before any merges; each plan obtained after a merge is represented by a data point in each of the figures. In the clinically most relevant part of the figures (treatment time under 5 minutes) the greedy approach provides a substantial improvement in the objective function and standard error of the PTV. The differences in the OAR sparing between the two VMERGE methods are minimal.	36
Figure 2.3	A comparison of treatment time versus plan quality trade-offs with the original VMERGE method, the greedy variant, and both methods with our SPG-regularization, for the prostate case. The two horizontal lines are objective function values of the optimal 9-beam and 20-beam step-and-shoot IMRT plans provided for reference. The SPG-regularization lowers the objective function values for both merging strategies, and greedy merging combined with SPG-regularization provides the most desirable plans.	38
Figure 2.4	A comparison of the two VMERGE methods, without SPG regularization, for the paraspinal case. In (a), we observe that for a fixed treatment time, the greedy VMERGE method achieves a lower objective function. This is especially pronounced in the most clinically relevant part of the graph, with treatment time under five minutes. (b) Mean dose to the spinal cord and kidneys, which are the two dose-limiting OARs.	39
Figure 2.5	A comparison of treatment time versus plan quality trade-offs with the original and greedy VMERGE methods, with and without regularization, for the paraspinal case. The horizontal lines are objective function values for 9-beam and 20-beam step-and-shoot IMRT plans. As in the prostate case, using greedy merging improves the objective function value for the same treatment times, and adding SPG-regularization lowers the objective function value even more substantially. A treatment plan with the same plan quality as the optimal 20-beam IMRT plan can be delivered in just over two minutes.	40
Figure 3.1	The cell survival curve given by the linear quadratic (LQ) model in (3.1). The solid line shows the estimation from the LQ model, and the data and error bars are from in-vitro experiments [18]. Reprinted with permission from Elsevier. . . .	43

Figure 3.2	Patient geometries for each of the five cases. The structures contoured on the CT scan are the liver (red), GTV (orange), chest wall (light green), large and small bowel (yellow), stomach (dark blue), esophagus (light blue), kidneys (purple), spinal cord (light brown), heart (magenta), and duodenum (dark green).	58
Figure 3.3	Dose distributions for Case 1, a large central lesion within the liver. (a) Physical dose distributions in each of the five fractions show that the nonuniformly fractionated treatment hypofractionates different parts of the tumor. (b) Total physical dose delivered throughout the nonuniformly fractionated treatment. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan, which is the uniform plan that is isoeffective in delivering the same BED as the nonuniformly fractionated plan. (e) The difference between the physical dose in the uniform plan and the DEQ5 for the nonuniform plan, or (c) minus (d). This shows that the spatiotemporal plans reduce dose in the healthy liver and in the entrance region of the those beams that expose the liver the most. All numerical quantities shown are in [Gy].	59
Figure 3.4	Rows (a), (b), and (c) are physical dose distributions from three additional locally optimal 5-fraction nonuniform treatments for Case 1, which is the same clinical liver case displayed in Figure 3.3. They are all locally optimal solutions of the model (3.4). The first five panels of each row are the dose distributions in the five nonuniform fractions, and the last panel on each row is the equivalent dose DEQ5. The solutions exhibit the same pattern: different subregions of the tumor receive a high single-fraction dose in different fractions. Note that the emergent “partitions” in the optimized plans are a result of the optimization. This pattern supports the rationale that the benefit of spatiotemporal fractionation is a result of hypofractionating parts of the tumor while maintaining a consistent low dose in the surrounding tissue. The difference in the hypofractionated regions in each solution also demonstrates that several qualitatively different locally optimal spatiotemporal treatments may exist for the same case. All numerical quantities shown are in [Gy].	60
Figure 3.5	Dose-volume histogram for (a) total physical dose and (b) DEQ5 for various structures in Case 1. The dashed lines are the values from the uniformly fractionated reference plan and the solid lines are from the spatiotemporal plan. These curves indicate that the spatiotemporal plan achieves an overall reduction in physical dose, as observed by the lines shifted to the left in (a). Additionally, in (b) we note that the spatiotemporal plan maintains DEQ5 in the tumor while reducing dose to the healthy liver tissue. To assess the conformity of the dose around the PTV, we included a curve for the voxels in a 1-cm margin around the PTV.	61

Figure 5.1	(a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 3 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.	93
Figure 5.2	Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 3, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.	94
Figure 5.3	A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 3 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.	95
Figure A.1	Dose distributions for Case 2, a small lesion within the liver. (a) Physical dose distributions in a 5-fraction nonuniformly fractionated treatment. Although the tumor is small, the nonuniform plan still hypofractionates different subregions of the target. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan. As in Case 1, the nonuniform spatiotemporal plan provides similar target coverage as the uniform reference plan, with a more conformal dose distribution. (e) The plot of (c) minus (d) shows that, as expected, the majority of the BED reduction is in the healthy liver and in the entrance regions of the beams exposing the liver the most. All numerical quantities shown are in Gy.	115
Figure A.2	Nonuniformly fractionated dose distributions for Case 3, which contains two lesions within the liver. (a) Physical dose distributions in each of the five fractions show that the nonuniformly fractionated treatment hypofractionates different parts of the tumor. (b) Total physical dose delivered throughout the nonuniformly fractionated treatment. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan, which is the uniform plan that is isoeffective in delivering the same BED as the nonuniformly fractionated plan. (e) The difference between the physical dose in the uniform plan and the DEQ5 for the nonuniform plan, or (c) minus (d). This shows that the benefit of nonuniform fractionation is mostly in reduced dose in the healthy liver and in the entrance region of the those beams that expose the liver the most. All numerical quantities shown are in Gy.	116

Figure A.3	Dose distributions for Case 4, a lesion abutting the chest wall. (a) Physical dose distributions for each fraction of a nonuniformly fractionated treatment. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniform reference plan. (d) DEQ5 of the nonuniformly fractionated plan. (e) Plot of (c) minus (d). All numerical quantities shown are in Gy.	117
Figure A.4	Dose distributions for Case 5, a lesion abutting the GI tract. (a) Physical dose distributions for each fraction of a nonuniformly fractionated treatment. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan. (e) Plot of (c) minus (d). All numerical quantities shown are in Gy.	118
Figure B.1	(a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 1 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.	120
Figure B.2	Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 1, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.	121
Figure B.3	A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 1 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.	121
Figure B.4	(a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 2 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.	122

Figure B.5	Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 2, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.	123
Figure B.6	A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 2 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.	123
Figure B.7	(a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 4 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.	124
Figure B.8	Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 4, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.	125
Figure B.9	A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 4 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.	125
Figure B.10	(a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 5 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.	126

- Figure B.11 Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 5, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together. 127
- Figure B.12 A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 5 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy. 127

Chapter 1

Introduction

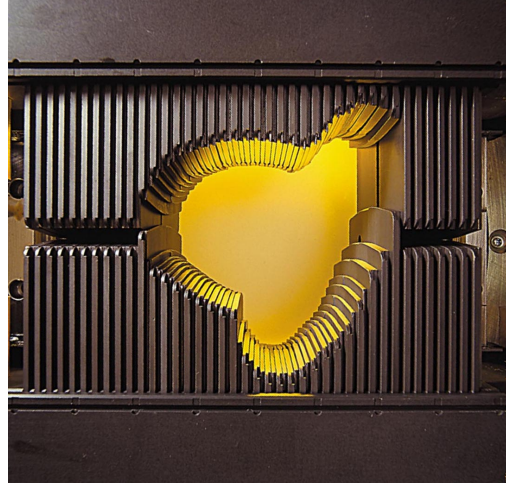
Radiotherapy is a widely used and highly effective modality to treat cancer, and it is a coalescence of the fields of physics, medicine, biology, and mathematics. Radiation was first used for therapeutic purposes at the turn of the twentieth century, and throughout its history, technological and scientific developments in each of these fields has propelled the progress of radiation for cancer treatment [23]. In 2015, more than half of all cancer patients were treated with some type of radiotherapy, either in conjunction with other treatment modalities such as surgery, chemotherapy, or immunotherapy, or as the primary treatment [47]. Despite the fact that it is so widely used, there remain many unanswered questions, and there is much potential for further progress in the field. In this thesis, we focus on mathematical optimization problems that arise in radiation therapy.

1.1 External beam radiotherapy

In external beam radiotherapy, ionizing radiation is generated outside of the patient and aimed at a tumor to damage malignant cells and prevent them from replicating. Typically, several radiation beams with non-lethal magnitudes are delivered from various angles around the patient. The cells in the intersection of the beams receive the highest amount of radiation, while the cells affected by a single beam receive a sub-lethal dose.



(a)



(b)

Figure 1.1: (a) The equipment setup for radiation treatment. The patient lies on the flat table, and the linear accelerator rotates around the patient on a gantry. (b) Close-up image of a multi-leaf collimator, which is positioned in front of the source of the radiation beam in the linear accelerator to control the beam's size and shape. Images courtesy of Varian Medical Systems, Inc., Palo Alto, CA. Copyright 2019. All rights reserved.

In photon radiotherapy, photon beams are generated by a linear accelerator, which is mounted on a rotating gantry. The patient is positioned on a table, and the gantry rotates around the patient to administer the radiation treatment (see Figure 1.1a). The shape of each radiation beam is controlled by a multileaf collimator (MLC), which is an apparatus that is positioned between the photon source and the patient. The MLC is made of pairs of “leaves”, or small blocks of a heavy metal that prevent radiation from passing through (see Figure 1.1b). The leaf positions are each controlled separately, and they affect the size and shape of the beam aperture, which in turn controls which cells are exposed to radiation.

Before a patient receives radiation treatment, a CT image of the treatment region is taken. On the CT scan, the structures in the patient are contoured, labeled, and classified into one of several types in order to prioritize them and specify how much radiation the structures should receive. Some healthy organs in the patient are called *organs at risk* (OARs), which must receive as little radiation as possible to avoid the side effects of radiation treatment. The observable tumor

of cancerous cells is labeled as the *gross tumor volume*, or GTV; these are cells that must be irradiated the most in order to have an effective treatment. In many radiation treatment plans, a small margin around the GTV is contoured and labeled as well. This region, called the *clinical target volume* (CTV), is contoured to include cancer cells that may be undetectable on the CT scan. Oftentimes an additional margin is added around the CTV to form a structure called the *planning target volume* (PTV). The PTV is included in treatment plans to be conservative for uncertainties in treatment delivery. Later in this thesis, (see Chapter 5) we will discuss the role of the PTV in treatment planning with uncertainty. Some clinical cases may contain additional targets, such as the lymph nodes for head and neck cases; these targets do not occur in the clinical cases used in this thesis, but they could be incorporated into the optimization models presented here in the same manner as the primary target structures. Any remaining tissue that is not part of the structures above is considered “Unclassified” (UNC).

1.2 Radiotherapy treatment modalities

The earliest and simplest form of radiotherapy is *3D conformal* radiotherapy, which delivers photon beams through an aperture that is the same shape as the beam’s-eye-view of the PTV. The development of the MLC allowed radiation beams to have varying intensities across the beam in a modality called *intensity-modulated radiotherapy* (IMRT). IMRT can either be delivered with beams that are fixed at certain gantry angles, or it can be delivered with *arc therapy*, where the beam is on while the gantry rotates around the patient. Less commonly, beams of heavy charged particles such as protons, helium ions, or carbon ions are used to damage malignant cells, but these require more complex and expensive equipment than photon radiotherapy.

1.2.1 Fixed beam intensity-modulated radiotherapy (IMRT)

A fixed beam IMRT treatment delivers photon beams from fixed gantry angles, but unlike the 3D conformal method it delivers modulated beams in which the intensity of the photons can be non-constant across the beam. This inhomogeneity is achieved by delivering multiple unmodulated

beams with different intensities from the same angle; the effects of the beams are superimposed, as illustrated in Figure 1.2. At a given gantry angle, a beam is delivered through some MLC leaf configuration. The beam is then turned off, the MLC leaves are rearranged, and the beam is turned back on *at the same gantry angle* to deliver another beam, possibly with a different intensity level. The total effect of this superposition is a beam with inhomogeneous intensity. Fixed beam IMRT is also sometimes referred to as “step-and-shoot” IMRT because the beam is on only when the gantry is stopped at a specified gantry angle. Typical clinical IMRT treatments use between five and nine distinct beam angles.

1.2.2 Photon arc therapy (VMAT)

Photon arc therapy, also known as intensity-modulated arc therapy (IMAT) or volumetric modulated arc therapy (VMAT), is a technique in which a beam is constantly on while the gantry is rotating around the patient. In addition, the MLC leaves are continuously moving as the gantry is rotating. VMAT is an attractive treatment modality because it uses the same equipment as fixed-beam IMRT, and it has the potential to greatly reduce treatment time. However, VMAT treatment planning is a large-scale, nonconvex optimization problem that must take additional parameters, such as dose rate and gantry speed, into account. Thus far, there is no universally accepted method for VMAT planning optimization among numerous approaches [91], and in this thesis we use a technique known as VMERGE [27].

1.2.3 Intensity-modulated proton therapy (IMPT)

Intensity-modulated proton therapy (IMPT) uses proton beams instead of photon beams to damage malignant cells. The optimization algorithms for proton and photon therapies are similar, but there are key differences in the absorption of protons versus photons in human tissue. The absorption of dose with each beam type is depicted in the depth-dose curve found in Figure 1.3a, which shows the amount of radiation absorbed in a tissue as a function of the tissue’s depth in the patient. With a proton beam, the dose absorption remains mostly constant until it peaks at a point called a Bragg

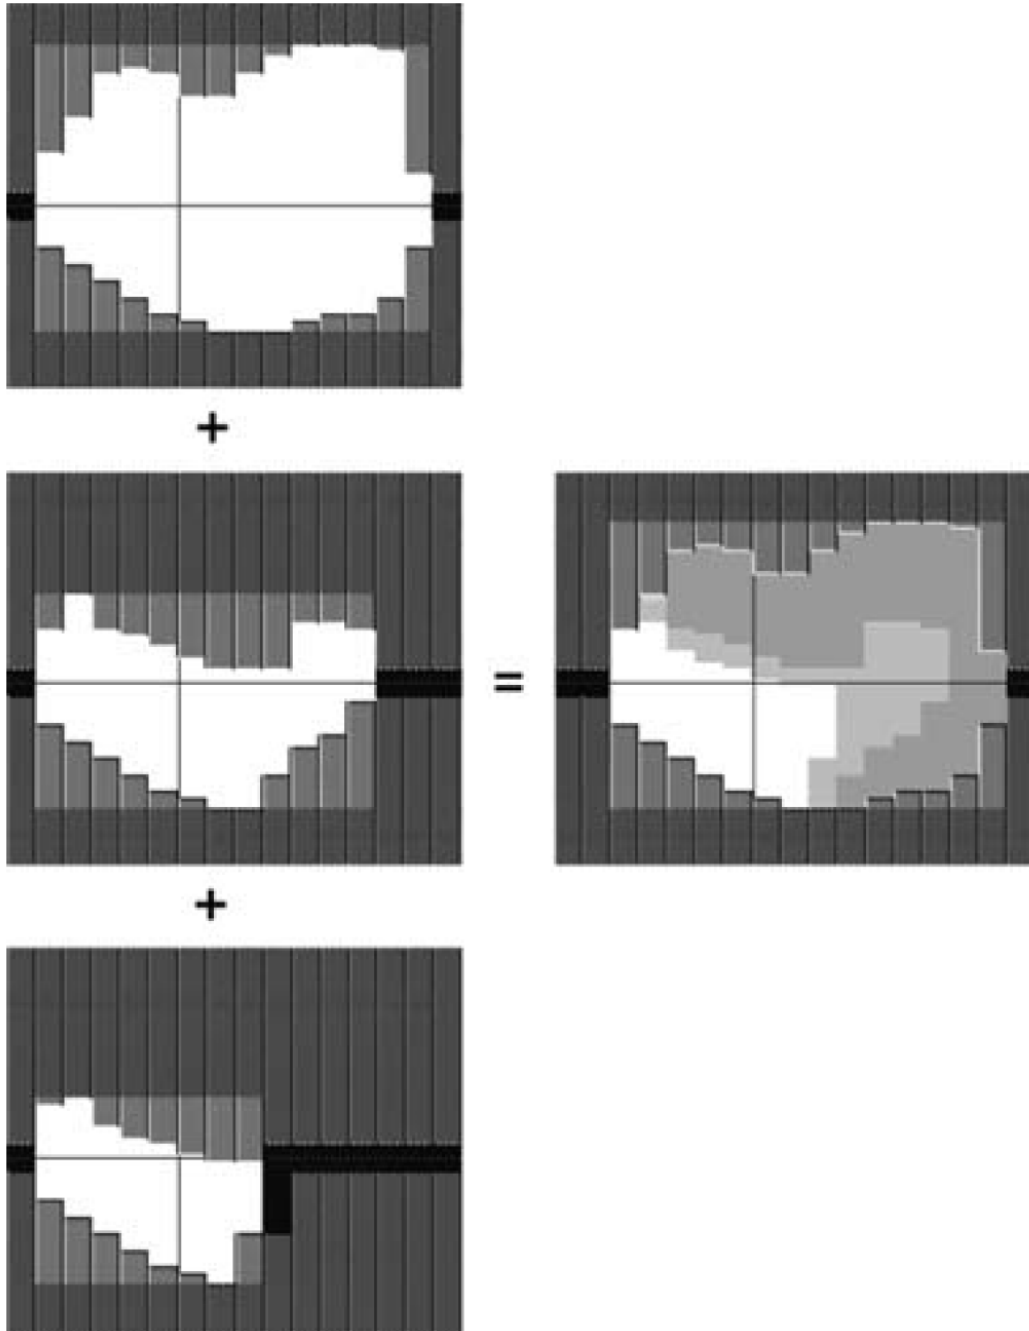


Figure 1.2: This graphic represents an example of a fixed-beam intensity-modulated radiotherapy (IMRT) treatment [75]. Three photon beams with different intensity levels are delivered through the three apertures on the left in succession from the same beam angle, and the resulting beam on the right has an inhomogeneous intensity that is the sum of the intensities in each of the beams on the left. Reprinted by permission from the publisher.

peak and then recedes quickly. In contrast, the photon depth-dose curve has a large peak at a more shallow depth and a much slower fall-off, meaning that cells in front of and behind the tumor will also be irradiated. A comparison of proton and photon dose distributions for the same patient can be found in Figure 1.3b.

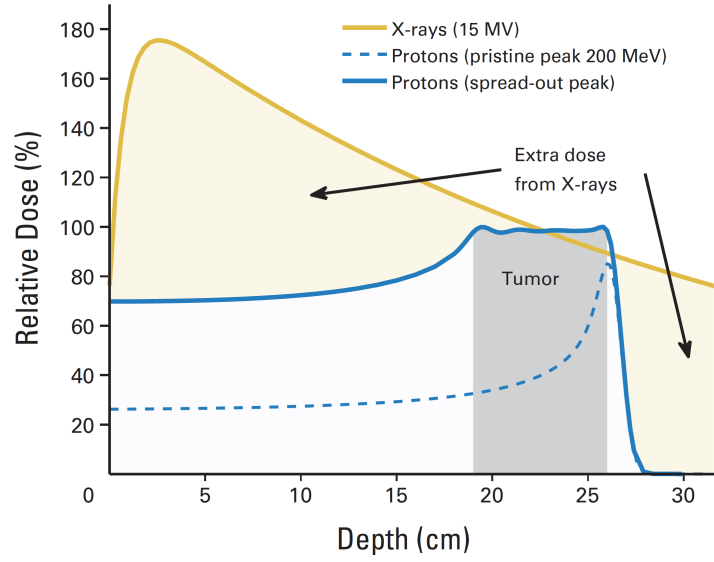
Currently, proton treatment facilities are far less prevalent than photon facilities because proton therapy requires specialized equipment, including particle accelerators such as cyclotrons or synchrotrons. As of October 2017, there were only 64 operational proton therapy facilities worldwide [69], whereas there are over 2,000 radiotherapy facilities with a photon linear accelerator in the United States alone [2].

1.3 Optimization of radiotherapy treatments

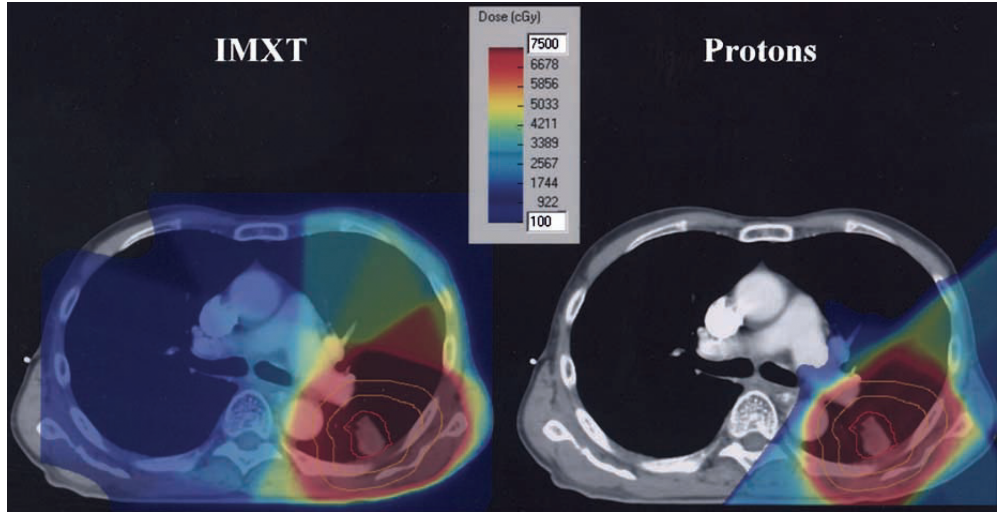
In IMRT treatment planning, the cross-section of a radiation beam is discretized into a rectangular grid of *beamlets* whose intensities are simultaneously optimized. The radiation to be delivered is represented by the *fluence maps*, which are the arrays of intensities assigned to each beamlet. Each beam angle is associated with a single fluence map.

There are two phases of optimization in radiotherapy treatment planning: the first is *fluence map optimization* (FMO), in which the optimal fluence maps are found that will yield the most effective treatment plan for target coverage and healthy tissue sparing. The next phase of the optimization is *leaf sequencing*, which finds the series of MLC leaf positions that will deliver a treatment plan as close as possible to the one that was optimized in FMO. For fixed-beam IMRT, there are standard leaf sequencing algorithms that can very closely replicate any fluence map. Thus, in this thesis we focus on the fluence map optimization phase of radiotherapy treatment planning.

At a very high level, the fluence map optimization problem is to select the vector of beamlet intensities that provides a treatment which meets prescription doses of radiation in the tumor volume(s) while delivering as little radiation to the healthy tissue as possible. In this section, we will introduce the conventional optimization models based on physical dose in the context of fixed-beam IMRT. Variations will be later introduced as needed.



(a)



(b)

Figure 1.3: (a) A comparison of the depth-dose curves for photon and proton beams [60]. Reprinted with permission. © 2014 American Society of Clinical Oncology. All rights reserved. (b) A comparison of the dose absorption within a patient using photon and proton beams [86]. (The acronym IMXT is intensity-modulated x-ray therapy, which is the same as IMRT.) Consistent with the depth-dose curves in (a), the dose distribution from the photon treatment delivers dose in front of and behind the tumor, while the proton beam avoids healthy tissues behind the tumor. Reprinted with permission from the publisher.

1.3.1 Clinical objectives

A multicriteria approach

In radiotherapy treatment planning, an oncologist prescribes some amount of radiation dose to the GTV and restricts the dose delivered to certain OARs. Each of these requirements represents a clinical objective, and typically there are many, oftentimes conflicting, clinical objectives for a single treatment plan. In this way, the optimization of treatment plans can be viewed as a multicriteria optimization (MCO) problem. In general, an MCO problem aims to simultaneously optimize multiple objective functions f_1, \dots, f_p , with $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for each i , over a feasibility region $\mathcal{X} \subset \mathbb{R}^n$ in a problem of the form

$$\min_{x \in \mathcal{X}} \{f_1(x), \dots, f_p(x)\}$$

[33, Ch. 2]. The notion of an optimal solution in MCO is fundamentally different from optimality for problems with a single objective function. A solution $\hat{x} \in \mathcal{X}$ is called *Pareto efficient* if there is no $x \in \mathcal{X}$ such that $f_k(x) \leq f_k(\hat{x})$ for $k = 1, \dots, p$ and $f_i(x) < f_i(\hat{x})$ for some $i \in \{1, \dots, k\}$ [33, Ch. 2]. In other words, a feasible solution \hat{x} is Pareto efficient if there is no other x in the feasible region that improves on one objective i without compromising another objective. The collection of all Pareto efficient points is called the *Pareto surface*, or *Pareto frontier*.

Often, choosing a solution of an MCO problem involves navigating the Pareto surface to explore the trade-offs between the objectives. There is much previous work on using MCO formulations of optimization problems in radiotherapy treatment planning to find treatment plans that balance the clinical objectives [11, 29, 78]. In many applications, a way to handle MCO problems is to combine the objectives into a single objective function by using the weighted sum

$$\sum_{i=1}^p w_i f_i(x)$$

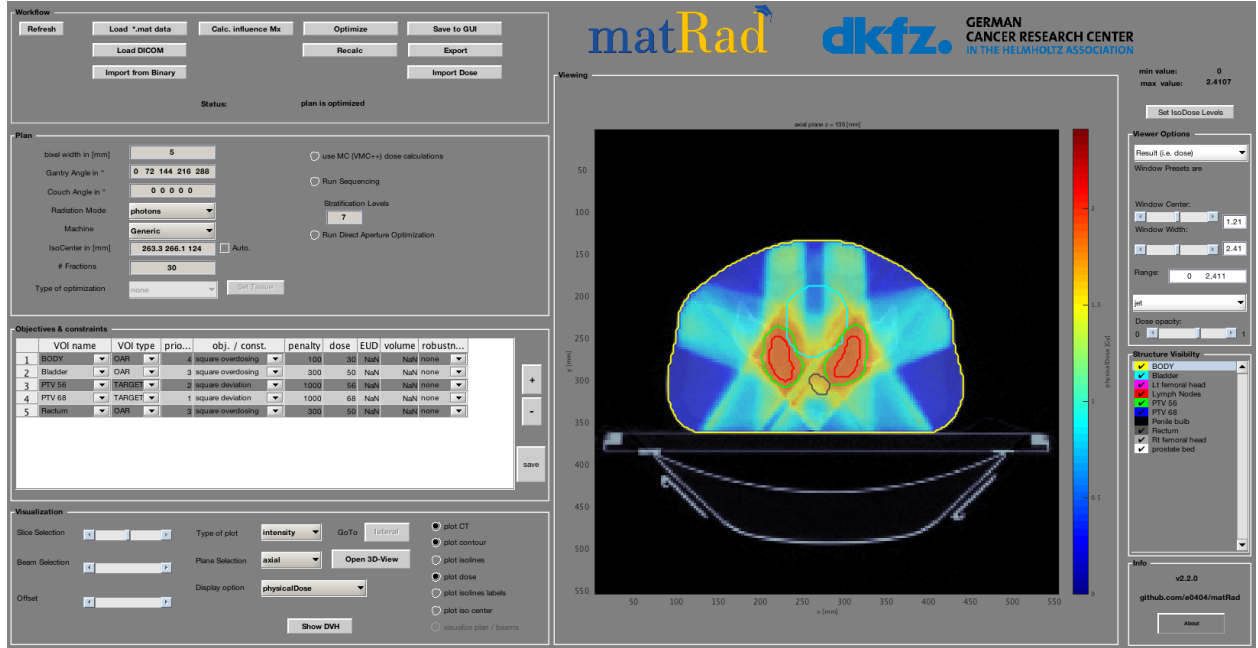


Figure 1.4: A screenshot of the open-source treatment planning system matRad [104]. Within the interface of a treatment planning system, a dosimetrist can complete all the steps of the treatment planning process, which includes specifying prescription values for each structure, calling the optimization solver, and evaluating treatment plan quality.

with $w_i \in \mathbb{R}_+$, for $i = 1, \dots, p$. The minimum of this sum is Pareto efficient [57]. Treatment planners use a treatment planning system (see Figure 1.4) to find solutions on the Pareto surface that balance the clinical objectives in a satisfactory way. We will also use a weighted sum of clinical objectives in the optimization problems throughout this thesis.

Quantifying the quality of a dose distribution

Metrics like the mean, maximum, or minimum dose in certain structures are utilized to evaluate how well a plan conforms to the clinical objectives. The metric for a particular organ depends on whether the organ can be described as a *serial* organ or a *parallel* organ, though not all organs fit neatly into one of these two categories. For serial organs, the functionality depends on the health of the entire organ. An example of a serial organ is the spinal cord; complications can occur if even a small part of the spinal cord is irradiated. Thus, the clinical objective for a serial organ is to limit

the maximum dose so that no single part of the structure suffers too much damage. In contrast to serial organs, parallel organs are able to sustain high amounts of damage in one part of the organ, and the undamaged area of the organ can function relatively normally. A mean dose objective for parallel organs reflects this property. An example of a parallel organ is the liver, which is the primary dose-limiting organ in the cases studied in Chapters 3 through 5 in this thesis. For organs that are not perfectly parallel or serial, the p -norm of the dose distribution (with $1 < p < \infty$) is often used to quantify the dose level in the organ. In radiotherapy nomenclature this is called *equivalent uniform dose* (EUD) [65, 88].

In addition to specifying goals for minimum, maximum, or mean dose values to various structures of the patient, clinicians can specify goals called *dose-volume constraints* such as “95% of the GTV must receive at least 50 Gy of dose”. While popular in the clinic, these dose-volume constraints are nonconvex and are therefore usually approximated in optimization models. A diagram called a *dose volume histogram* (DVH) is used to visually represent these objectives when evaluating plan quality. An example of a dose volume histogram can be found in Figure 1.5. Each structure is represented by a different curve on the DVH, and the curve indicates the percentiles of physical dose within the structure. For example, in Figure 1.5, approximately 10% of the bladder receives 40 Gy or more of physical dose. DVH plots are widely used in the medical community to evaluate the quality of a treatment plan.

1.3.2 A mathematical model of fluence map optimization

To formulate a mathematical model, the patient is discretized into a 3-dimensional grid of *voxels*, and the structures of the patient, such as organs, target volumes, or unclassified tissue, are represented as sets of voxels. The beam aperture is discretized into a grid of *beamlets*, and the arrays of radiation intensities assigned to each beamlet are called *fluence maps*, which are the decision variables in the optimization problem. Figure 1.6 is a pictorial representation of the optimization problem.

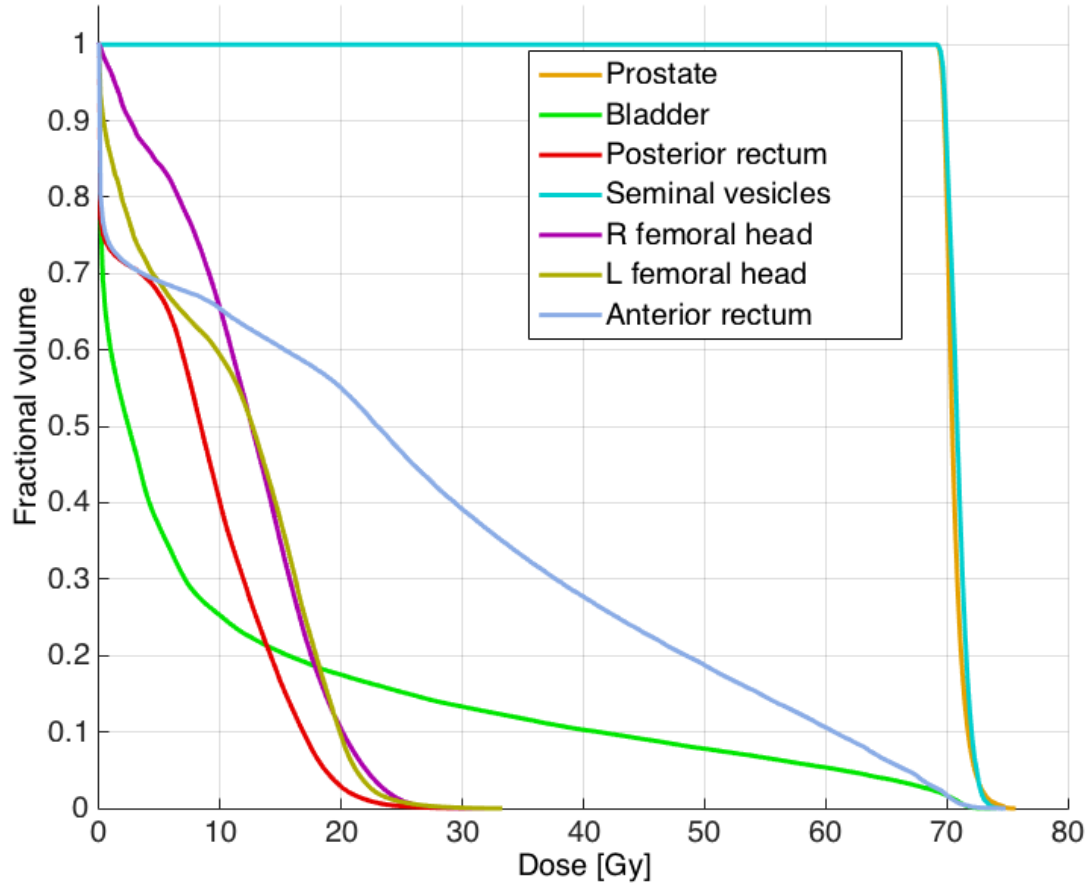


Figure 1.5: An example of a dose volume histogram (DVH) used to evaluate the quality of a treatment plan for a prostate case. Each curve represents a different structure within the patient, and for a fixed amount of dose, the DVH displays the fraction of the structure which receives at least that amount of dose. The prostate and seminal vesicles are the target volumes with a prescription dose of 70 Gy, and the remaining structures are OARs.

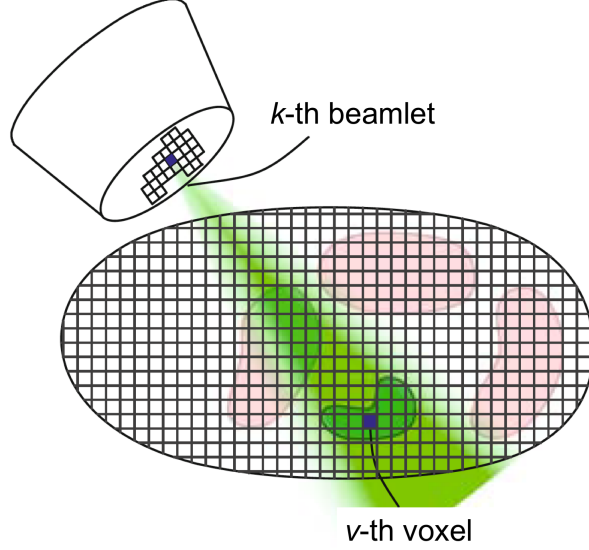


Figure 1.6: A pictorial representation of the setup of an IMRT problem [75]. The oval figure at the bottom represents the patient, which is divided into a grid of voxels. In the gantry head, the beam aperture is divided into a grid of beamlets, and the beamlet intensities are the decision variables in the fluence map optimization problem. Reprinted by permission from the publisher.

Let $B \in \mathbb{N}$ be the total number of beamlets for all of the beams, and let V be the set of all voxels in the patient, with the total number of voxels denoted by $|V|$. The fluence maps are denoted by the vector $x = (x_1, \dots, x_B)$, where each element x_k is the intensity of the k -th beamlet measured in *monitor units* (MU). The physical doses delivered to each voxel of the patient are denoted by the vector $d = (d_1, \dots, d_{|V|})$, where each d_v is the dose delivered to voxel v , measured in the physical unit *gray* (symbol: Gy). The physical dose that is absorbed by each voxel of the patient is a linear function of the beamlet intensities delivered by the linear accelerator; the linear model is given by the expression $d = Dx$ with the *dose-influence matrix* $D \in \mathbb{R}^{|V| \times B}$. The (i, j) -th entry in the dose matrix is the dose delivered to voxel i by a unit intensity in beamlet j . In this work, a dose-influence matrix for each patient is computed with an existing dose calculation algorithm and is considered to be a fixed part of the patient data. In a typical clinical case with fewer than 10 beams, the number of beamlets B is on the order of thousands, and the number of voxels $|V|$ is in the tens of millions.

With this notation established, the fluence map optimization problem is as follows:

$$\begin{aligned}
& \min_{x \in \mathbb{R}^B, d \in \mathbb{R}^{|V|}} F(d) \\
& \text{s.t.} \quad d = Dx \\
& \quad \quad x \geq 0.
\end{aligned} \tag{1.1}$$

The optimal fluence map minimizes a function of the physical dose d , subject to only two constraints. First, the dose-influence relationship from the linear model must be maintained, and second, the beamlet intensities must be nonnegative (because a negative amount of radiation cannot be delivered). In the implementation, the dose vector d is eliminated using the matrix multiplication Dx , but in this thesis the dose vectors d are included in the optimization models for clarity.

Let I represent the index set of clinical objectives in the treatment plan prescription. In this thesis, prescriptions are defined in terms of three types of objectives: prescription upper bounds and lower bounds for individual voxels and mean dose upper bounds for entire structures. Three types of penalty functions reflect these prescription types. Let I^+ and I^- be the index sets of objectives that penalize dose above and below the prescribed thresholds, respectively, and let I^{m+} be the index set of objectives penalizing mean dose over the prescription limit. The index set I is the union of the sets I^+ , I^- , and I^{m+} . Let V_i denote the set of voxels that are associated with clinical objective i .

It is possible that a given voxel could be subjected to multiple penalties simultaneously; a voxel could have a prescription upper bound and a prescription lower bound, and the voxel could be a member of a structure that has mean overdose penalties. If objective i imposes voxelwise prescription limits, let d_{iv}^{hi} and d_{iv}^{lo} denote the prescription upper and lower bounds, respectively, for voxel v . Similarly, if objective i is a bound for the mean dose of some structure, let m_i^{hi} denote the upper bound for the mean dose. Deviations from these prescription values can be penalized using the sum of one-sided deviations squared, where the one-sided deviation is represented by the positive part function given by $(y)_+ = \max\{0, y\}$. The piecewise quadratic penalty functions of the

physical dose F_i are defined as

$$F_i(d) = \begin{cases} \sum_{v \in V_i} (d_v - d_{iv}^{\text{hi}})_+^2 & \text{if } i \in I^+, \\ \sum_{v \in V_i} (d_{iv}^{\text{lo}} - d_v)_+^2 & \text{if } i \in I^-, \\ \left(\frac{1}{|V_i|} \sum_{v \in V_i} d_v - m_i^{\text{hi}} \right)_+^2 & \text{if } i \in I^{m+}. \end{cases} \quad (1.2)$$

Typically, a treatment plan uses the same threshold d_i^{hi} or d_i^{lo} in every voxel v of the same structure V_i , but we opted for the above, more general, formulation to allow for different thresholds for voxels within the same structure. For example, in later chapters we use a *conformity objective* in which the prescription threshold for a voxel around the PTV is a function of the voxel's distance from the PTV.

The objective function in the optimization problem (1.1) is a weighted sum of the individual penalty functions:

$$F(d) = \sum_{i \in I} w_i F_i(d). \quad (1.3)$$

The weights $w_i \in \mathbb{R}_+$ reflect the relative importance of each clinical objective, and they are chosen by the treatment planner, either directly or indirectly via Pareto surface navigation [26, 28, 38, 62]. Since the F_i are all convex, every Pareto optimal solution can be obtained by minimizing (1.3) with some set of positive weights [33, Ch. 3].

The objective function (1.3) is a (once) continuously differentiable convex function, and when used in the optimization problem (1.1) together with the substitution $d = Dx$ in the objective function, we have an essentially unconstrained problem except for nonnegativity constraints of the beamlet intensities. A globally optimal solution can be found using standard convex optimization solvers. The IMPT optimization model is mathematically analogous to the IMRT model presented in this section, but VMAT plans are optimized with different large-scale nonconvex models to

simultaneously optimize the leaf trajectories, beam intensities, and gantry rotation in a VMAT plan.

1.4 Fractionation of radiotherapy treatments

Radiotherapy treatments are typically *fractionated*, which means the total dose is divided among treatments that are delivered over a series of days or weeks. Each of the individual treatments is called a *fraction*. Fractionated treatments allow for the preservation of healthy tissue because healthy cells are able to recover from sub-lethal radiation damage. Traditionally, radiation treatments are *uniformly fractionated*, meaning that the total dose is divided evenly into equal doses on each treatment day. A typical uniform fractionation scheme has between 5 and 30 fractions of treatment, and the same dose distribution is delivered in every fraction.

Fractionated treatments are ideal for healthy tissue sparing, but to deliver lethal damage to tumor cells, an ideal treatment would deliver a high amount of dose in a small number of fractions (this is called a *hypofractionated* treatment). Thus, there is a trade-off between fractionating the treatment for healthy tissue sparing and hypofractionating the treatment for effective tumor coverage.

1.4.1 Spatiotemporal fractionation

Unlike uniform fractionation that delivers the same dose distribution in each fraction, *spatiotemporal fractionation*, (also called *nonuniform fractionation*) is a relatively novel approach that allows dose distributions to vary between the fractions. (The terms *nonuniform fractionation* and *spatiotemporal fractionation* are used synonymously throughout this thesis.) Spatiotemporally fractionated treatments are able to at least approximately achieve both of the goals of hypofractionating parts of the tumor while fractionating dose to the healthy tissue. The concept of spatiotemporal fractionation is quite recent, and it was first proposed in the context of IMPT [98]. Spatiotemporal fractionation schemes have not been implemented in the clinic thus far because their benefit has not been properly assessed, particularly in the context of photon therapy. This thesis contributes to the

growing body of literature that demonstrates that spatiotemporal fractionation schemes are able to substantially improve photon treatment plan quality over conventionally fractionated treatments.

1.5 Uncertainty in radiotherapy treatments

Precision in the delivery of a radiotherapy treatment plan is imperative because an incorrect delivery can lead to severe degradation of the treatment. If a treatment is not delivered in accordance with how it was planned, tumor cells could be spared from radiation, compromising the effectiveness of the treatment, and healthy tissues could be subjected to radiation levels that could cause potentially fatal side effects. In this section, we briefly introduce sources of uncertainty that may affect the accurate delivery of radiotherapy treatments and summarize a mathematical approach to mitigate uncertainties.

1.5.1 Sources of uncertainty in radiotherapy

There are two broad classes of uncertainty in radiation therapy: systematic errors and random errors. *Systematic* errors occur consistently throughout the treatment, while *random* errors exhibit no regular pattern. Potential sources of error in treatment planning and delivery include particle range uncertainty in IMPT [70, 71, 93], variability in target volume delineation [5, 107], or errors in patient setup during treatment delivery [4, 14, 39, 71]. Radiotherapy treatments could also be compromised because of anatomy changes in the patient [100] or organ motion, including respiratory motion [14, 15, 20].

The severity of each type of uncertainty depends on the treatment site (in the body) and treatment modality; for example, respiratory motion would be a more pressing concern in a lung cancer treatment than in a prostate cancer treatment. In this thesis, we focus on random patient setup errors that could occur during the treatment. In conventional treatment planning, setup error is usually handled using margins around the target, such as the PTV [99], which ensure that the GTV is not too severely underdosed if the patient position is shifted slightly. We will discuss later

in Chapter 5 that the concern of adequate GTV coverage in the presence of random patient setup uncertainty is even more pronounced for spatiotemporal plans.

1.5.2 Optimization under uncertainty

Stochastic optimization is one approach to handle uncertainties in optimization problems. Stochastic optimization models contain functions of a random vector ω that is assumed to have a probability distribution P supported on $\Omega \subset \mathbb{R}^s$, where s is the dimension of the uncertainty. An example of a simple stochastic program is

$$\min_{x \in X} \mathbb{E}_P [F(x, \omega)],$$

where x is a deterministic decision vector on $X \subset \mathbb{R}^n$ and $F : X \times \Omega \rightarrow \mathbb{R}$ [80, Ch. 1]. In addition, stochastic programs may have probabilistic constraints such as

$$P\{g(x, \omega) \leq 0\} \geq 1 - \alpha,$$

for some $\alpha \in [0, 1]$ and $g : X \times \Omega \rightarrow \mathbb{R}$, which requires that the inequality $g(x, \omega) \leq 0$ must hold at the $1 - \alpha$ probability level [80, Ch. 4]. In Chapter 5, we formulate stochastic optimization models to optimize spatiotemporal treatments that are robust against patient setup uncertainty.

1.6 The contribution of this thesis

In Chapter 2, we propose two modifications for an arc therapy planning algorithm called VMERGE with the goal of improving the achieved trade-off between treatment time and plan quality. We compare this trade-off for both the original algorithm and the algorithm with our proposed modifications using two clinical cases, and we observe noticeably improved treatment plans for the same treatment times over what can be obtained using the original VMERGE method. For one of the proposed modifications, we present an algorithm that exploits the problem structure to optimize treatments efficiently.

In Chapter 3, we present optimization models for spatiotemporal fractionation that are based on a quadratic phenomenological model called the biologically effective dose (BED) model. This formulation leads to a computationally challenging nonconvex optimization problem. For five clinical liver cases, we demonstrate that spatiotemporal fractionation schemes improve plan quality over conventionally fractionated treatments in photon arc therapy.

In Chapter 4, we derive a bound on the maximum achievable benefit from spatiotemporal fractionation. Both Chapter 3 and prior works on spatiotemporal fractionation in the context of IMPT demonstrate that locally optimal solutions are high-quality treatment plans, but it has been unclear if better optimization algorithms may yield substantially better solutions. In this chapter we formulate a convex semidefinite programming relaxation of the nonconvex problem in Chapter 3 whose solution provides a rigorous upper bound on the maximum achievable benefit of spatiotemporal fractionation.

In Chapter 5, we address the concern of patient setup uncertainty in dose delivery of spatiotemporally fractionated treatments. Patient setup uncertainty is incorporated into the optimization models using expected penalty functions, and a discrete probability distribution is defined to characterize the random setup error occurring in each fraction. We demonstrate that in the presence of patient setup uncertainty, spatiotemporally fractionated treatment plans exhibit robust tumor coverage and retain some of the benefit in sparing healthy tissues.

Some of the material in this thesis has been published in peer-reviewed journals. Parts of Chapter 2 were published as a Technical Note in *Medical Physics* [40]. The contents of chapters 3 and 4 were published in *Physics in Medicine & Biology* [41]. A manuscript based on Chapter 5 is in preparation.

Chapter 2

Improving photon arc therapy delivery

We revisit the VMERGE treatment planning algorithm by Craft *et al.* [27] for arc therapy planning and propose two changes to the method that are aimed at improving the achieved trade-off between treatment time and plan quality at little additional planning time cost, while retaining other desirable properties of the original algorithm. The VMERGE algorithm is a two-phase approach, where a high-quality initial plan is determined in the first phase; the delivery of this plan is usually very time-consuming, and the plan may even be undeliverable due to machine constraints such as minimum gantry speed. The initial plan is made deliverable in a practical amount of time in the second phase. We propose two changes to this approach, which affect both phases. We implement a “sum-of-positive-gradients” (SPG) smoothing in the first phase that constrains the treatment time of the initial plan, and in the second phase we use a greedy algorithm to merge maps in a way that keeps the dose degradation low. The SPG-constrained treatment can be optimized with an efficient algorithm that exploits the structure of the problem.

2.1 Introduction

Despite the fact that photon arc therapy (VMAT) has found widespread clinical application, as mentioned in Section 1.2.2, the optimization of VMAT plans remains a challenging large-scale nonconvex problem, with a variety of competing approaches to its (incomplete) solution [73, 91]. The VMERGE algorithm, proposed by Craft *et al.* [27], is one of these approaches. VMERGE has several desirable properties: it is very fast, it is compatible with multicriteria optimization (MCO), and it allows the planner to select a desired trade-off between plan quality and delivery time. In the rest of this section we summarize the original VMERGE radiotherapy treatment planning method. Our suggested changes to the method follow in Section 2.2. In Section 2.3 we describe an efficient algorithm to solve the SPG-constrained problem, and in Section 2.4 we present our numerical results using a prostate and a paraspinal case.

The VMERGE algorithm

As described in Section 1.2.1, static beam IMRT treatments deliver intensity-modulated beams from fixed gantry angles, and in our implementation of fluence map optimization (FMO), the beamlet intensities for each beam angle are optimized using the piecewise quadratic objective function (1.3) to penalize under- and over-dosing of various organs. Let b denote the number of beam angles in the static-beam IMRT treatment. In Section 1.3.2, we used the notation $x \in \mathbb{R}^B$ to denote the fluence maps, where B is the total number of beamlet intensities to optimize. In this chapter, because the spatial location of each beamlet in the field is important for leaf sequencing, we represent the k -th fluence map by using the matrix $x^k \in \mathbb{R}^{r \times q}$, for $k = 1, \dots, b$, where the grid of beamlet intensities has r rows and q columns. Relating this back to the notation from Chapter 1, the total number of beamlet intensities to optimize over all beams is $b \times r \times q = B$. Throughout this work we assume every beam has a rectangular grid of beamlets of the same size.

The first step of the VMERGE algorithm is solving FMO on an angular grid of fine (say, 2-degree) resolution. The attained 180 fluence maps represent an approximation of a perfect treat-

ment that could be delivered to the patient by a single-arc VMAT, but this 180-beam fluence map is not feasible to deliver to a patient. First, it is not possible to achieve arbitrarily high intensity modulation during a 2-degree rotation of the gantry; the maximum MLC leaf speed and the minimum gantry speed are mechanical limitations on some machines that prevent a highly-modulated field from being delivered in a small 2-degree arc sector. Even if such intensity modulation within a 2-degree angle is possible, the amount of time to deliver this treatment would be prohibitive; for example, in the clinical cases presented later in this chapter, the delivery time for treatments with 180 fluence maps is over 19 minutes, whereas a static-beam IMRT plan usually takes less than 10 minutes and a VMAT plan takes less than 3 minutes.

To make the 180-beam deliverable, the second phase of the VMERGE algorithm iteratively merges similar adjacent fluence maps and their corresponding gantry arc sectors. With each merge, the number of fluence maps decreases by one. After merging, the resulting treatment plan can be accurately delivered with a full VMAT arc in a reasonable amount of time. The merging of adjacent fluence maps is dictated by the similarity between maps, which is measured by the Frobenius norm of the arc-length-weighted fluence maps. More precisely, the similarity score δ for two maps x^1 and x^2 that are to be delivered in adjacent arc sectors with arc lengths θ_1 and θ_2 , respectively, is given by

$$\delta(x^1, x^2) = (\theta_1 + \theta_2) \left(\sum_{i=1}^p \sum_{j=1}^q \left(\frac{x_{ij}^1}{\theta_1} - \frac{x_{ij}^2}{\theta_2} \right)^2 \right)^{1/2}, \quad (2.1)$$

where i, j indicate the row and column indices of the fluence maps [27]. The pair of adjacent fluence maps with the lowest δ is added and delivered over the union of their individual arc sectors; this will reduce the treatment time, but it could result in a degradation of plan quality. The merging of adjacent maps continues until the dose degradation has surpassed a chosen threshold. The stopping criterion can be specified in terms of the violation of the original clinical goals, which is likely a better practice than using a threshold based on the less directly interpretable optimization objective function used in the FMO.

2.2 Two improvements to the VMERGE algorithm

2.2.1 A greedy approach to merging

Our first change to the VMERGE algorithm is implementing a greedy strategy in choosing the pair of maps to merge. In the original VMERGE algorithm, the pair of maps to merge was determined by the similarity score δ defined in (2.1), but simply merging adjacent maps that are similar beamletwise does not necessarily make the merged map easier to deliver over the combined arc sector. Because of the way the MLC leaves sweep across the field, it may be more beneficial to merge adjacent maps that are not necessarily similar to each other, but whose combination results in a dose distribution that is similar to the ideal distribution when delivered over the combined arc sector. To quantify this idea, for each pair of adjacent fluence maps, we compute the dose degradation that would result from the merging of those maps, which is measured using the same objective function that is used in the initial FMO step of the algorithm. In each iteration, we merge the pair of maps whose merging results in the smallest dose degradation.

As in the original VMERGE, two maps are merged by adding the fluence of the corresponding beamlets and delivering the new map over the union of the individual arc sectors. Merging maps with the greedy variant is computationally more expensive than merging based on the similarity score δ , but only marginally so. Suppose that a certain number of merges have already been performed, and the dose distribution d delivered by the plan obtained after the last merging step has already been computed. Then, in order to evaluate the dose degradation resulting from merging two given adjacent maps, we only need to compute the dose corresponding to the two fluence maps to be merged and the dose corresponding to the merged map. Since each map can only be merged with one of two adjacent maps, the total computational work needed to evaluate the dose degradation resulting from the merging of all adjacent pairs is equivalent to two complete dose calculations for a 180-beam plan. With a precomputed dose-influence matrix D , the dose d corresponding to the plan x is simply $d = Dx$, and the total computational work required in each iteration is equivalent to two matrix-vector multiplications involving the matrix D .

2.2.2 Regularization and treatment time constraints in fluence map optimization

Our second change in the VMERGE implementation affects the treatment plan with 180 fluence maps that is computed in the first step of the algorithm. As mentioned above, this “ideal” treatment plan is undeliverable because it contains an arbitrarily high amount of intensity modulation within a single fluence map. If adjacent beamlets have very different beamlet intensities, then it is more time-consuming to deliver that fluence map because the MLC leaves must move across the field in small increments. It would be more desirable for the 180-map plan to have some amount of “smoothness” in the fluence maps, or similarities between adjacent beamlet intensities, so that the MLC leaves can move across the field more quickly and the treatment time can be reduced.

Quantifying the delivery time for a fluence map is a challenging combinatorial problem with no known efficient solution [91]. However, for a *sliding window* delivery (also known as a *sweep* in dynamic multileaf collimation), the time required to deliver a fluence map can be computed exactly, assuming that the dose rate is constant [87]. During this mode of delivery, all collimator leaves start on one side of the field and move unidirectionally across the field. Leaves may only turn back once all leaves have reached the other side of the field. The amount of time it takes for a pair of collimator leaves to move across the field is given by the expression $\frac{w}{\ell}$, where w is the physical width of the field (which equals the beamlet width times q) and ℓ is the maximum leaf speed. The additional time required to deliver the modulated dose for the row $x_{i,\cdot}^k$ of a fluence map is given by the *sum-of-positive-gradients* (or *SPG*) function, expressed as

$$\text{SPG}(x_{i,\cdot}^k) = x_{i,1}^k + \sum_{j=1}^{q-1} \max \left\{ x_{i,j+1}^k - x_{i,j}^k, 0 \right\}. \quad (2.2)$$

Combining the dose modulation time with the time for the leaves to sweep across the field, the delivery time for the i -th row of fluence map k is given by

$$t(x_{i,\cdot}^k) = \frac{w}{\ell} + \frac{1}{\alpha} \left(\text{SPG}(x_{i,\cdot}^k) \right), \quad (2.3)$$

where α is the constant dose rate.

It can be shown using elementary arithmetic that SPG can also be expressed using the formula

$$\text{SPG}(x_{i,\cdot}^k) = \sum_{j=1}^q x_{ij}^k - \sum_{j=1}^{q-1} \min \left\{ x_{i,j}^k, x_{i,j+1}^k \right\}, \quad (2.4)$$

which reveals that the SPG is the same whether the collimator leaves move right-to-left or left-to-right. The time to deliver the k -th fluence map x^k with sliding window delivery is the maximum of the delivery times of each collimator leaf pair (that is, of each row of x^k). Using the notation $x = (x^1, \dots, x^b) \in \mathbb{R}^B$ for the entire fluence vector of a treatment plan with b beams, the total delivery time for treatment plan x is

$$T(x) = \sum_{k=1}^b \max_{i=1,\dots,r} t(x_{i,\cdot}^k). \quad (2.5)$$

To achieve the goal of having “smoother” fluence maps when computing the ideal 180-beam IMRT plan, we explicitly include an upper bound constraint on the treatment time, assuming sliding window delivery. We refer to the resulting plan as the *regularized solution*. To obtain the regularized solution, we aim to find the optimal set of nonnegative beamlet weights x that minimize the physical dose-based piecewise quadratic penalty function, subject to an upper bound U on the treatment time. Formally, the optimization problem to obtain the regularized solution is

$$\begin{aligned} \min_x \quad & F(x) \\ \text{s.t.} \quad & T(x) \leq U \\ & x \geq 0, \end{aligned} \quad (2.6)$$

where the objective function $F(x)$ is the convex penalty function defined in Section 1.3.2 and $T(x)$ is the total delivery time defined in (2.5).

Note that $\text{SPG}(\cdot)$ is a convex piecewise affine function, and as a result, so is $T(\cdot)$; thus, we can write the constraint $T(x) \leq U$ as a collection of linear constraints by introducing the auxiliary

variables m_k for $\max_{i=1,\dots,r} t(x_{i,\cdot}^k)$ and the auxiliary variables $p_{i,j}^k$ for $\max\{x_{i,j+1}^k - x_{i,j}^k, 0\}$ in (2.2). The constraints $T(x) \leq U$ and $x \geq 0$ are satisfied if and only if there exist $m_k \in \mathbb{R}$ and $p_{i,j}^k \in \mathbb{R}$ (for $i=1, \dots, r$, $j=1, \dots, (q-1)$, $k=1, \dots, b$) that satisfy the system of constraints

$$\begin{aligned}
& \sum_{k=1}^b m_k \leq U \\
& \frac{w}{\ell} + \frac{1}{\alpha} \left(x_{i,1}^k + \sum_{j=1}^{q-1} p_{i,j}^k \right) \leq m_k \quad \forall k = 1, \dots, b, \quad \forall i = 1, \dots, r \\
& x_{i,j+1}^k - x_{i,j}^k - p_{i,j}^k \leq 0 \quad \forall k = 1, \dots, b, \quad \forall i = 1, \dots, r, \quad \forall j = 1, \dots, q-1 \\
& p_{i,j}^k \geq 0 \quad \forall k = 1, \dots, b, \quad \forall i = 1, \dots, r, \quad \forall j = 1, \dots, q-1 \\
& x_{i,j}^k \geq 0 \quad \forall k = 1, \dots, b, \quad \forall i = 1, \dots, r, \quad \forall j = 1, \dots, q.
\end{aligned} \tag{2.7}$$

In the clinical prostate case presented later in this chapter with 180 beams, this formulation contains over 72,000 variables and 109,000 constraints. The resulting optimization problem can be quite challenging to solve within practical time and memory limits, but in the next section we describe an algorithm to solve the optimization problem efficiently.

2.3 Optimizing a treatment plan with a constraint on the treatment time

To obtain the regularized solution from the optimization problem (2.6), we use an in-house implementation of the Frank–Wolfe algorithm (also known as the conditional gradient algorithm), which in general minimizes a convex function g over a convex domain \mathcal{D} [37, 48]. At an iterate x_t in the Frank–Wolfe algorithm, the next iterate is given by

$$x_{t+1} = (1 - \lambda_t)x_t + \lambda_t s, \tag{2.8}$$

where the step size $\lambda_t \in (0, 1]$ is given by $\lambda_t = \frac{2}{t+2}$ and s is the solution of the optimization problem

$$\min_{s \in \mathcal{D}} \nabla g(x_t)^\top s. \quad (2.9)$$

In other words, at every iteration of the algorithm, the first-order Taylor approximation of the objective function is minimized over the feasible region, and then a step is made in the direction of the minimizer. Variations of the Frank–Wolfe algorithm use different step sizes to promote sparsity of the solutions [22] or accelerate the convergence of the algorithm [44].

The Frank–Wolfe algorithm is especially useful if each instance of the inner problem (2.9) is simple to solve, e.g., when the feasible region is a norm ball or simplex [48]. Because the feasible region of (2.6) is a polyhedron, an interior point method [49] could be used to minimize the linear objective function over the feasible region (2.7). For this linear program, the number of variables and the number of constraints are both $\Theta(B)$. (Recall that $B = brq$ is the total number of beamlets, which for the clinical cases presented later in this chapter is in the tens of thousands.) An interior point method requires $\mathcal{O}(\sqrt{B})$ iterations, and in each iteration, one or more linear systems of size $\Theta(B)$ are solved. Thus, an interior point method could solve (2.9) in $\mathcal{O}(B^{3.5})$ time, requiring $\mathcal{O}(B^2)$ space for the linear system solves in each iteration. We will show in the remainder of this section that we can solve the inner problem (2.9) with substantially smaller time and space requirements.

First, we will establish an equivalent formulation of the optimization problem to illuminate its structure. Consider the polyhedron

$$P(\beta) := \left\{ x \in \mathbb{R}_+^{r \times q} \mid \text{SPG}(x_{i,\cdot}) \leq \beta \ \forall i = 1, \dots, r \right\}. \quad (2.10)$$

Using \hat{x} to denote the current iterate in the Frank–Wolfe algorithm and defining $\tilde{U} = \alpha(U - \frac{w}{\ell})$ and $\tilde{m} = \alpha(m - \frac{w}{\ell})$, the problem (2.9) is the linear program

$$\begin{aligned}
& \min_{\substack{x=(x^1, \dots, x^b) \\ m_1, \dots, m_b}} \quad \nabla F(\hat{x})^\top x \\
& \text{s.t.} \quad \sum_{k=1}^b \tilde{m}_k \leq \tilde{U} \\
& \quad \quad x^k \in P(\tilde{m}_k) \quad \forall k = 1, \dots, b,
\end{aligned} \tag{2.11}$$

which must be solved in every iteration of the algorithm.

In the following lemma, we note a useful property of the SPG function:

Lemma 2.3.1. *Let $x = (x_1, \dots, x_q) \in \mathbb{R}^q$, let $a \in \mathbb{R}$, and let $b \in \mathbb{R}$. Then $\text{SPG}(ax + b\mathbf{1}) = a\text{SPG}(x) + b$, where $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^q$.*

Proof. We have that

$$\begin{aligned}
\text{SPG}(ax + b\mathbf{1}) &= ax_1 + b + ((ax_2 + b) - (ax_1 + b))_+ + \dots + ((ax_q + b) - (ax_{q-1} + b))_+ \\
&= ax_1 + b + (ax_2 - ax_1)_+ + \dots + (ax_q - ax_{q-1})_+ \\
&= a(x_1 + (x_2 - x_1)_+ + \dots + (x_q - x_{q-1})_+) + b \\
&= a\text{SPG}(x) + b.
\end{aligned} \quad \square$$

From the homogeneity of the SPG function established in this lemma (with $b = 0$), an optimal solution of (2.11) is a scaled version of the problem

$$\begin{aligned}
& \min_{\substack{x=(x^1, \dots, x^b) \\ m_1, \dots, m_b}} \quad \nabla F(\hat{x})^\top x \\
& \text{s.t.} \quad \sum_{k=1}^b \tilde{m}_k \leq 1 \\
& \quad \quad x^k \in P(\tilde{m}_k) \quad \forall k = 1, \dots, b.
\end{aligned} \tag{2.12}$$

That is, if x^* is a solution to (2.12), then $\tilde{U}x^*$ is a solution to (2.11). For cleaner notation, throughout the remainder of this section we will use the formulation (2.12) for the inner problem of the Frank–Wolfe algorithm.

2.3.1 Decomposition of the inner problem

If it were not for the constraint $\sum_{k=1}^b \tilde{m}_k \leq 1$, the problem (2.12) could be decomposed into b subproblems with $P(\tilde{m}_k)$ as the feasible region, and each subproblem could be solved separately. While this trivial decomposition of the problem is not possible, we shall see in Theorem 2.3.2 that in an optimal solution of the inner problem (2.12), $\tilde{m}_k = 1$ for at most one index k and $\tilde{m}_k = 0$ for all other k . In the radiotherapy application, this means in the solution to the inner problem of the Frank–Wolfe algorithm, one beam is turned on and uses the total treatment time, while all the other beams are off, delivering no fluence. We will use the notation $c_k \in \mathbb{R}^{r \times q}$ for the elements of the gradient of the objective function corresponding to beam k , so that $\nabla F(\hat{x}) = c = (c_1, \dots, c_b)$. We prove in the following theorem that at least one of the optimal solutions of (2.12) has the structure described above:

Theorem 2.3.2. *For the optimization problem (2.12), there always exists an optimal solution x^* with $\max_{i=1, \dots, r} \text{SPG}(x_{i,\cdot}^{k*}) = 1$ for at most one index k^* and $\max_{i=1, \dots, r} \text{SPG}(x_{i,\cdot}^k) = 0$ for all other $k \neq k^*$.*

Proof. Suppose that $x^* = (x^1, \dots, x^b)$ is an optimal solution of (2.12). If $x^* = \mathbf{0}$, then the theorem holds. If x^* is not the zero vector, then let $\lambda_k := \max_{i=1, \dots, r} \text{SPG}(x_{i,\cdot}^k)$. Then $0 < \sum_{k=1}^b \lambda_k \leq 1$. In the case where $\sum_{k=1}^b \lambda_k < 1$, the solution can be scaled up so that $\sum_{k=1}^b \lambda_k = 1$ and the objective value is no worse than $c^\top x^*$. Suppose that at the optimal solution $x^* = (x^1, \dots, x^b)$, $\lambda_k > 0$ for at least two indices k , and without loss of generality assume $\lambda_1 > 0$ and $\lambda_2 > 0$. Then $c_1^\top x^1 \leq 0$ and $c_2^\top x^2 \leq 0$, because otherwise changing x^1 or x^2 to the zero vector would yield a better solution than x^* . Again without loss of generality, assume

$$\frac{c_1^\top x^1}{\lambda_1} \leq \frac{c_2^\top x^2}{\lambda_2}. \quad (2.13)$$

Now consider the vector $\bar{x} = \left(\frac{x^1(\lambda_1+\lambda_2)}{\lambda_1}, 0, x^3, \dots, x^b\right)$. With some algebraic manipulations we have

$$\begin{aligned}
c^\top \bar{x} &= \frac{(\lambda_1 + \lambda_2)c_1^\top x^1}{\lambda_1} + c_3^\top x^3 + \dots + c_b^\top x^b \\
&= c_1^\top x^1 + \frac{\lambda_2 c_1^\top x^1}{\lambda_1} + c_3^\top x^3 + \dots + c_b^\top x^b \\
&\stackrel{(2.13)}{\leq} c_1^\top x^1 + c_2^\top x^2 + c_3^\top x^3 + \dots + c_b^\top x^b \\
&= c^\top x^*.
\end{aligned}$$

Thus, the objective value of \bar{x} is no larger than the objective value of the optimal solution x^* . If the strict inequality holds, there is a contradiction because the feasible solution \bar{x} has a lower objective value than the vector x^* that was assumed to be optimal.

If the equality $c^\top \bar{x} = c^\top x^*$ holds, then \bar{x} is another feasible solution that has the same objective value as the optimal solution x^* and one fewer nonzero λ_k . Continuing with the remaining indices with $\lambda_k > 0$ in the same fashion as x^1 and x^2 , we can follow a similar procedure until we arrive at a contradiction or a solution where $\lambda_k > 0$ for exactly one k . \square

From the theorem above, we know that $\tilde{m}_k = 1$ for at most one value of k and $\tilde{m}_k = 0$ for all other k . If the index of the beam with $\tilde{m}_k = 1$ were known *a priori*, then solving the inner problem (2.12) would reduce to finding

$$(x^k)^* = \arg \min_{x \in P(1)} c_k^\top x \tag{2.14}$$

for a single value of k . Since this index is unknown, we obtain $(x^k)^*$ for every beam k and find a beam index \hat{k} with the lowest objective value. Then, the optimal solution of the inner problem (2.12) is $x^* = (0, \dots, (x^{\hat{k}})^*, \dots, 0)$.

In subsequent sections we will show that a subproblem of the form (2.14) can be solved in $\mathcal{O}(rq)$ time. Since the solution of (2.12) is obtained by solving (2.14) b times, the inner problem of the

Frank–Wolfe algorithm is solved in $\mathcal{O}(brq) = \mathcal{O}(B)$ time, as opposed to interior point methods that can solve the inner problem in $\mathcal{O}(B^{3.5})$ time.

2.3.2 The subproblem for each beamlet row

We can further decompose the problem by observing that the polyhedron $P(1)$ is the Cartesian product $Q \times Q \times \cdots \times Q$ with

$$Q := \{x \in \mathbb{R}_+^q \mid \text{SPG}(x) \leq 1\}. \quad (2.15)$$

Each $x^k \in P(1)$ has the form $x^k = (x_{1,\cdot}^k, \dots, x_{r,\cdot}^k)$ with each $x_{i,\cdot}^k \in Q$. The polyhedron Q is a bounded polyhedron (a *polytope*) because each component of x must be in the interval $[0, 1]$ to satisfy $\text{SPG}(x) \leq 1$. Using the notation c_k^i to denote the components of the objective function $\nabla F(\hat{x})$ corresponding to $x_{i,\cdot}^k$, we solve the subproblem (2.14) by computing

$$x_{i,\cdot}^k = \arg \min_{x \in Q} (c_k^i)^\top x \quad (2.16)$$

r times, once for each value of i . In the next section we will show that this reduced problem for a single beamlet row can be solved in $\mathcal{O}(q)$ time. Thus, the subproblem (2.14) for a single beam can be solved in $\mathcal{O}(rq)$ time.

We now provide the vertex representation of the polytope Q , which will aid in solving the reduced subproblem (2.16). Consider the set

$$\Xi := \{\xi \in \{0, 1\}^q \mid \exists i, j \in \{1, 2, \dots, q\} \text{ s.t. } \xi_k = 1 \text{ iff } i \leq k \leq j\} \cup \{\mathbf{0}\}.$$

The set Ξ is the set of all vectors in $\{0, 1\}^q$ with a contiguous set of nonzero elements. In the next two theorems, we will prove that Ξ is equal to the set of vertices of the polytope Q . We will first show that the set Ξ is a subset of the vertices of Q .

Theorem 2.3.3. *Every element in the set Ξ is a vertex of the polytope Q .*

Proof. Let $\xi \in \Xi$. A vector ξ is a vertex of the polytope Q if there exists a $c \in \mathbb{R}^q$ such that $c^\top \xi < c^\top y \forall y \in Q, y \neq \xi$ [8]. Consider the vector $c \in \{-1, 1\}^q$ defined by

$$c_k = \begin{cases} 1, & \text{if } \xi_k = 0 \\ -1, & \text{if } \xi_k = 1. \end{cases}$$

Let $y \in Q, y \neq \xi$, and since $y \in Q$, we know $y_i \in [0, 1] \forall i = 1, \dots, q$. If $\xi = \mathbf{0}$, then $c^\top y > 0 = c^\top \xi$. If $\xi \neq \mathbf{0}$, let i and j be the indices such that $\xi_k = 1$ iff $i \leq k \leq j$. Then $c^\top y > -(j - i + 1) = c^\top \xi$. Thus, ξ is a vertex of the polytope Q . \square

Now that we have established that every element in the set Ξ is a vertex of Q , we will also show that every vertex of Q is contained in Ξ . We will use the fact that $x \in Q$ is a vertex of Q if and only if it is an extreme point of Q , that is, a vector $x \in Q$ for which there are no other vectors $y, z \in Q, y \neq x$ and $z \neq x$, such that $x = \lambda y + (1 - \lambda)z$ for a scalar $\lambda \in (0, 1)$ [8]. We will show that every element in $Q \setminus \Xi$ can be written as a convex combination of two other elements in Q .

Theorem 2.3.4. *The set Ξ contains every vertex of the polytope Q .*

Proof. Let $x = (x_1, \dots, x_q) \in Q \setminus \Xi$, and let x_i be the smallest nonzero element of x . (The vector x has at least one nonzero element because $x \notin \Xi$.) Then $0 < x_i < 1$ because if $x_i = 1$ then $x = \mathbf{1} \in \Xi$, and if $x_i > 1$ then $x \notin Q$. Note that $x = x_i \mathbf{1} + (1 - x_i)z$ with $z = \left(\frac{x_1 - x_i}{1 - x_i}, \dots, \frac{x_q - x_i}{1 - x_i}\right) \in [0, 1]^q$. Since $x \in Q \setminus \Xi$, we know that $x \neq \mathbf{1}$, and since $x_i > 0 = z_i$, we also know that $x \neq z$. Further, applying Lemma 2.3.1,

$$\begin{aligned} \text{SPG}(z) &= \frac{1}{(1 - x_i)} \text{SPG}(x) - \frac{x_i}{(1 - x_i)} \\ &\leq \frac{1}{(1 - x_i)} - \frac{x_i}{(1 - x_i)} \\ &= 1, \end{aligned}$$

so $z \in Q$. Thus, x is a convex combination of two vectors (that are not the same as x) in Q . \square

From Theorems 2.3.3 and 2.3.4 we have that Ξ is exactly equal to the set of vertices of the reduced problem (2.16), so an optimal solution of the subproblem can be found in Ξ . The characterization of the vertices of Q with the set Ξ will allow the reduced problem for each beamlet row to be solved efficiently. One could solve the subproblem (2.16) in $\mathcal{O}(q^3)$ time by simply enumerating through all of the $\mathcal{O}(q^2)$ elements of Ξ and evaluating the objective function with $\mathcal{O}(q)$ arithmetic operations. Then the subproblem (2.12) could be solved in $\mathcal{O}(brq^3)$ operations; already, that is an improvement over an interior point method, which could solve (2.12) in $\mathcal{O}((brq)^{3.5})$ time. In the next section, we will demonstrate an even larger improvement with a linear-time algorithm to solve the reduced subproblem for a single beamlet row.

2.3.3 Linear-time algorithm to solve the reduced subproblem

We aim to find the vertex in Ξ with the lowest objective value. As described in the previous section, the vertices of the subproblem are vectors in $\{0, 1\}^q$ that have a contiguous set of nonzero elements. Thus, solving (2.16) is equivalent to finding the indices m and n , with $1 \leq m \leq n \leq q$, such that $\sum_{j=m}^n - (c_k^i)_j$ is maximized. This problem is known as the maximum subarray sum problem.

The maximum subarray sum problem can be solved with a dynamic programming algorithm known as Kadane's algorithm [43]. At index j , let c be the maximum sum for the subarray ending at position j , and let s be the maximum subarray sum obtained so far up to position j . The maximum subarray is extended if the sum of the subarray ending at position j plus x_{j+1} is greater than x_{j+1} . Otherwise, the maximum subarray stays the same and the iterations continue. Formally, Kadane's algorithm is given by

```

 $c \leftarrow \max\{0, x_1\}$ 
 $s \leftarrow \max\{0, x_1\}$ 
for  $j = 2 : \text{length}(x)$  do
     $c \leftarrow \max\{c + x_j, x_j\}$ 
     $s \leftarrow \max\{c, s\}$ .
end for
```

The indices m and n for the beginning and end of the subarray with the maximum sum can also be maintained throughout the iterations. Then the optimal solution of (2.16) is the vector x^* defined as

$$x_j^* = \begin{cases} 1, & \text{if } m \leq j \leq n \\ 0, & \text{if } j < m \text{ or } j > n. \end{cases}$$

This algorithm can be used to solve the reduced subproblem (2.16) in linear ($\mathcal{O}(q)$) time. Then, since the subproblem is solved r times for each of the b beams, the problem (2.12) in each iteration of the Frank–Wolfe algorithm can be solved in $\mathcal{O}(brq)$ time.

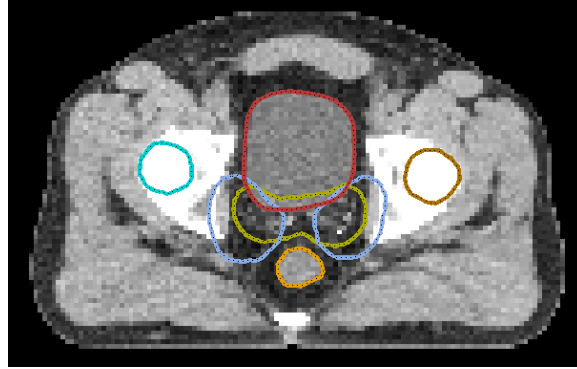
2.4 Numerical experiments with clinical cases

We examined the effect of the proposed changes to the VMERGE planning method using two clinical cases: the first is the prostate case from the CORT dataset [25], and the second is a paraspinal tumor with a more challenging geometry (see Figure 2.1). The prostate case is a standard one with a small number of organs-at-risk (bladder, rectum, penile bulb, and femoral heads) and two targets: a geometric expansion of the prostate is prescribed 68 Gy in 30 fractions, and an expansion around the prostate and the lymph nodes is prescribed 56 Gy in 30 fractions.

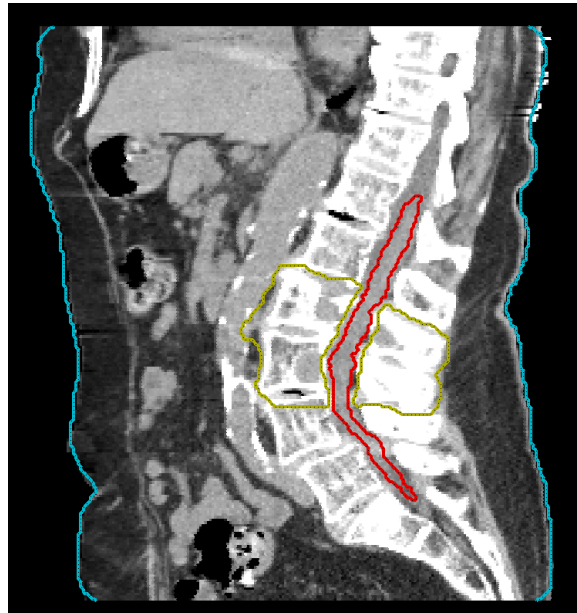
The paraspinal case has a single target volume that surrounds the spinal cord, which is the main dose-limiting organ. The secondary dose-limiting organs in this case are the kidneys. The prescribed dose to the PTV is 66 Gy in 33 fractions. For the OARs, the maximum dose in the center of the spinal cord is limited to 50 Gy, and the mean dose in the kidneys is constrained to be under 25 Gy.

2.4.1 Experimental setup

Single-arc VMAT plans were computed using different variants of VMERGE. We compared the two merging strategies, the one based on the original similarity score (2.1) and one based on the greedy strategy from Section 2.2.1. Both of these were employed with and without the SPG-based



(a)



(b)

Figure 2.1: Representative slices of CT images for the clinical cases in this chapter. (a) The prostate case from the CORT dataset. The visible contoured OARs in this slice are the bladder (red), left and right femoral heads (light blue and brown), and the rectum (orange). The two PTVs are contoured in blue and yellow. (b) The paraspinal case. The PTV (contoured in yellow) is wrapped around the spinal cord, which is contoured in red.

regularization from Section 2.2.2. All of the plans were optimized assuming sliding window MLC leaf trajectories with a maximum leaf speed of 6 cm/s and a maximum dose rate of 600 MU/min.

Dose deposition coefficients were computed for 180 equispaced 6 MV beams with 1×1 cm beamlet grid resolution using the dose calculation algorithm implemented in CERR version 4.0b4 [31]. The maximum field size was adapted to each case; the field width was determined by the smallest rectangular field that covers the beams-eye-view of the target from each of the 180 beam angles. For the prostate case, this rectangular field is 13×16 cm, for the paraspinal case it is 9×12 cm.

The different variants of the VMERGE algorithm were started from an optimal 180-beam step-and-shoot IMRT plan (with and without the SPG constraint), and the merges were continued until only 10 merged fluence maps remained. Each plan following a merge represents a different trade-off between delivery time and treatment quality. In both of the clinical cases, the initial 180-sector plan required far too much delivery time, and the quality of the 10-sweep plan was far too poor to be useful. The overall plan quality was measured by the FMO objective function value. For visualization and evaluation purposes, two additional metrics were also used [24]:

- the target coverage was quantified by the standard error of the (voxelwise) PTV dose,

$$\text{std_error} = \left(\frac{1}{P} \sum_{i=1}^P (M - d_i)^2 \right)^{1/2} \quad (2.17)$$

where P is the number of voxels in the PTV, and M is the prescription dose;

- organ sparing was quantified by the mean dose to the respective organs.

2.4.2 Results

Results from the prostate case

Figure 2.2 compares the merging strategies of the original VMERGE method and the greedy variant. The first panel shows the difference in terms of our primary plan quality metric, the FMO objective

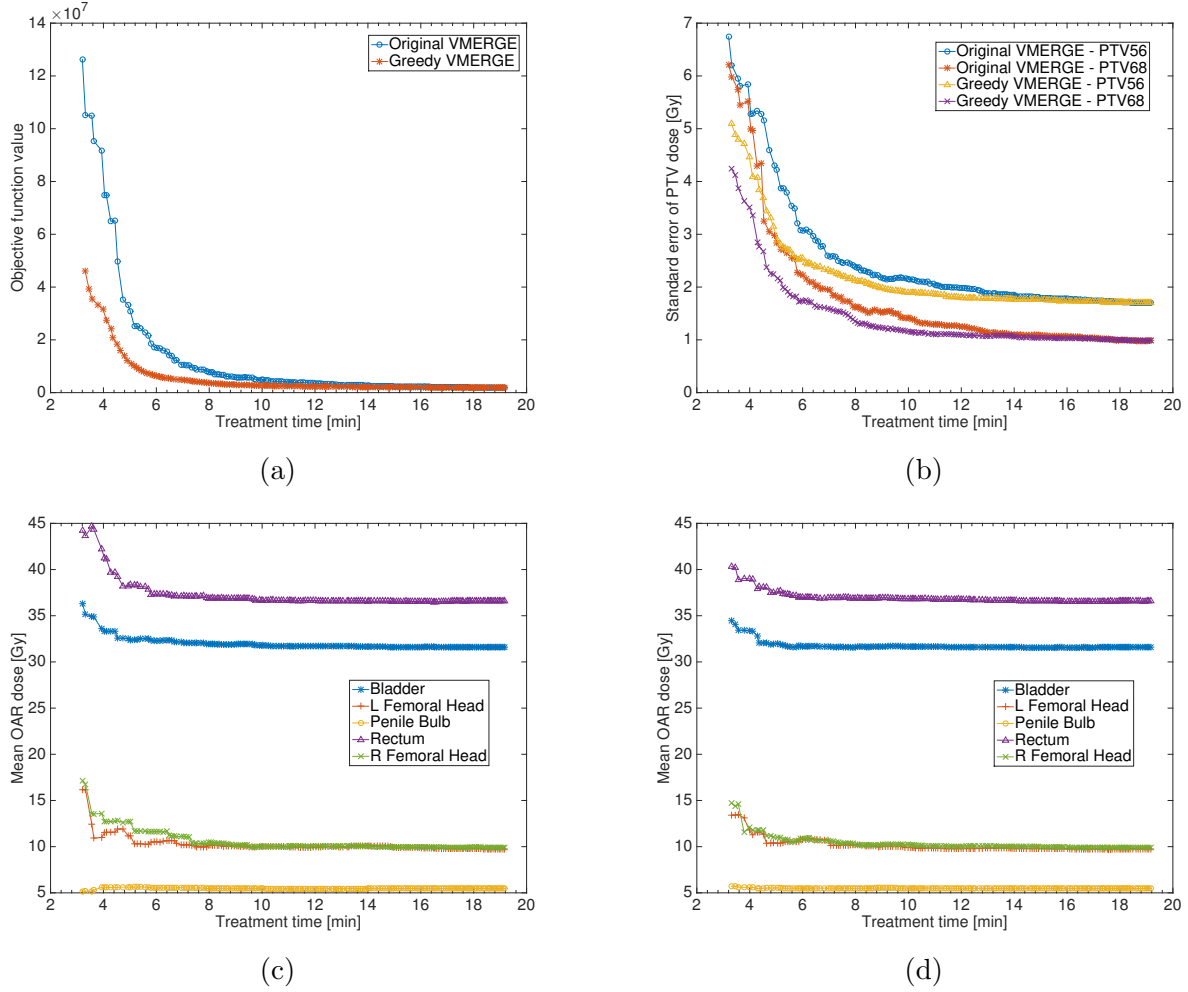


Figure 2.2: Comparison of the original VMERGE method and the greedy variant on the prostate case, without SPG-regularization. Panel (a): objective function value versus treatment time. Panel (b): standard error of PTV dose (Eq. (2.17)) versus treatment time for the two PTVs. Panel (c): mean dose to the OARs with the original VMERGE method. Panel (d): mean dose to the OARs with the greedy variant. The data points with the highest treatment time represent the 180-beam IMRT plan before any merges; each plan obtained after a merge is represented by a data point in each of the figures. In the clinically most relevant part of the figures (treatment time under 5 minutes) the greedy approach provides a substantial improvement in the objective function and standard error of the PTV. The differences in the OAR sparing between the two VMERGE methods are minimal.

function value. The remaining panels show the standard error of the PTV dose and the mean dose to each OAR separately. In the most relevant section of graph, corresponding to 2–6 minutes of treatment time, for two plans that have the same treatment quality, the plan obtained with the greedy merging strategy has around 1.5 minutes lower treatment time. The greedy variant also yields lower voxelwise standard error for both of the PTVs in this case.

Figure 2.3 shows the effect of starting VMERGE from a regularized SPG-constrained solution. As outlined in Section 2.2.2, these plans were obtained by starting with a 180-beam FMO solution that was optimized with the additional constraint that the delivery time of the plan should be under a specific time. For this case the time constraint was chosen to be 540 seconds; it takes 480 seconds to sweep across the 16 cm wide field 180 times without delivering any dose (with the collimator leaves traveling at their maximum speed), and we allowed one additional minute to deliver radiation and modulate the radiation field. To better gauge the quality of these treatment plans, we marked in the figure the objective function values of two fixed-beam IMRT plans (with no constraints on treatment time or modulation of the fluence), using 9 and 20 equispaced beams. Fixed beam IMRT treatments typically have between 5 and 9 beams; hence, the 9-beam plan is a high-quality plan that could realistically be delivered in the clinic, while the 20-beam plan is better than typical clinical treatments. The fixed beam IMRT plans were obtained by optimizing the same FMO objective function that is used throughout the study to measure treatment quality.

The treatment plans benefit considerably from both modifications to the VMERGE method; in fact, the effect of the regularization appears to be even greater than the effect of the greedy merging strategy. Combining regularization and greedy merging, our modified VMERGE method provided a treatment comparable to 9-beam step-and-shoot IMRT with an estimated treatment time under 3 minutes.

Results from the paraspinal case

In Figure 2.4, we observe that the paraspinal case has the same qualitative results as the prostate case. For a fixed treatment time, the greedy VMERGE algorithm yields a solution with a lower

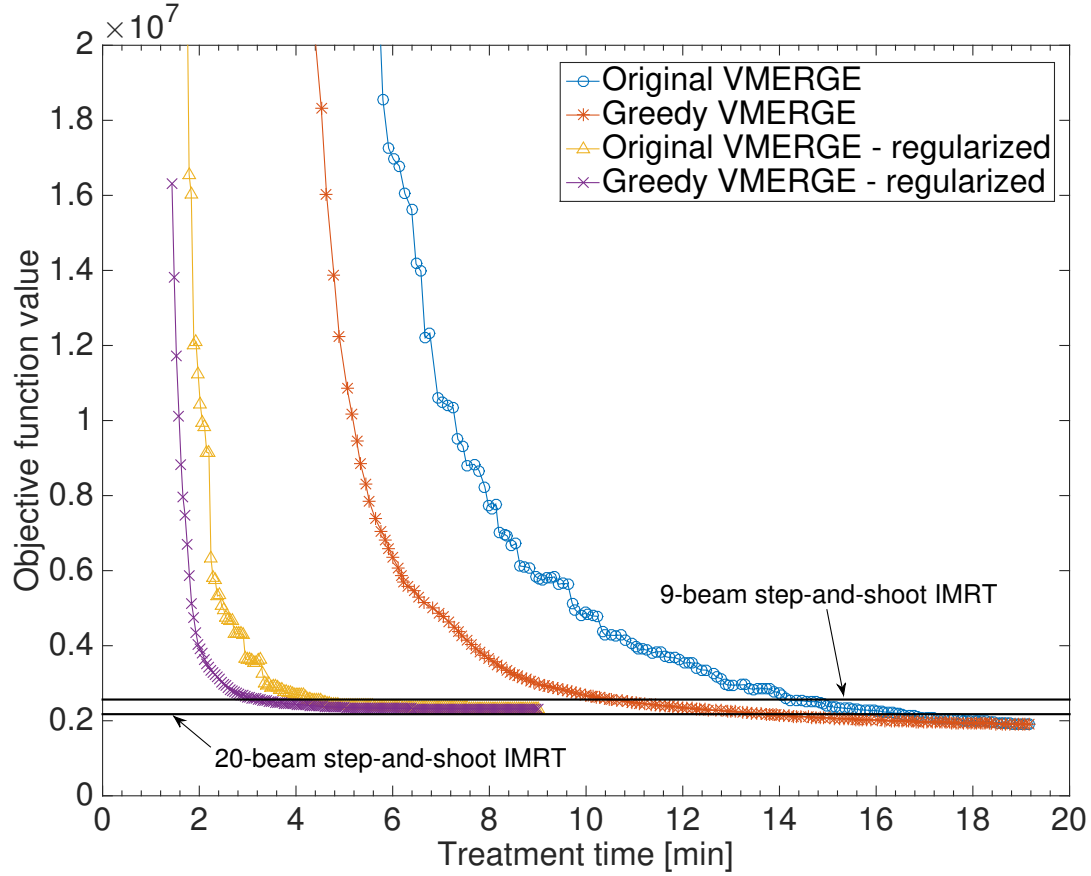
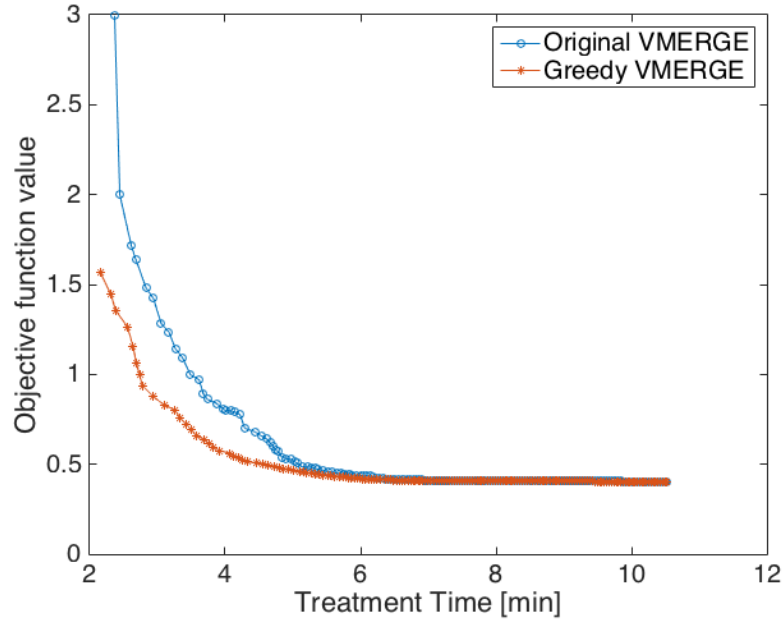
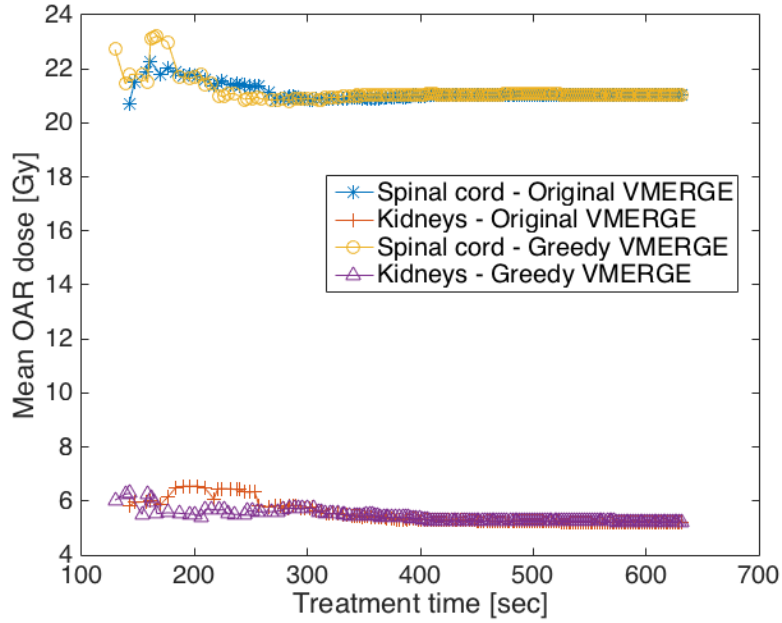


Figure 2.3: A comparison of treatment time versus plan quality trade-offs with the original VMERGE method, the greedy variant, and both methods with our SPG-regularization, for the prostate case. The two horizontal lines are objective function values of the optimal 9-beam and 20-beam step-and-shoot IMRT plans provided for reference. The SPG-regularization lowers the objective function values for both merging strategies, and greedy merging combined with SPG-regularization provides the most desirable plans.



(a)



(b)

Figure 2.4: A comparison of the two VMERGE methods, without SPG regularization, for the paraspinal case. In (a), we observe that for a fixed treatment time, the greedy VMERGE method achieves a lower objective function. This is especially pronounced in the most clinically relevant part of the graph, with treatment time under five minutes. (b) Mean dose to the spinal cord and kidneys, which are the two dose-limiting OARs.

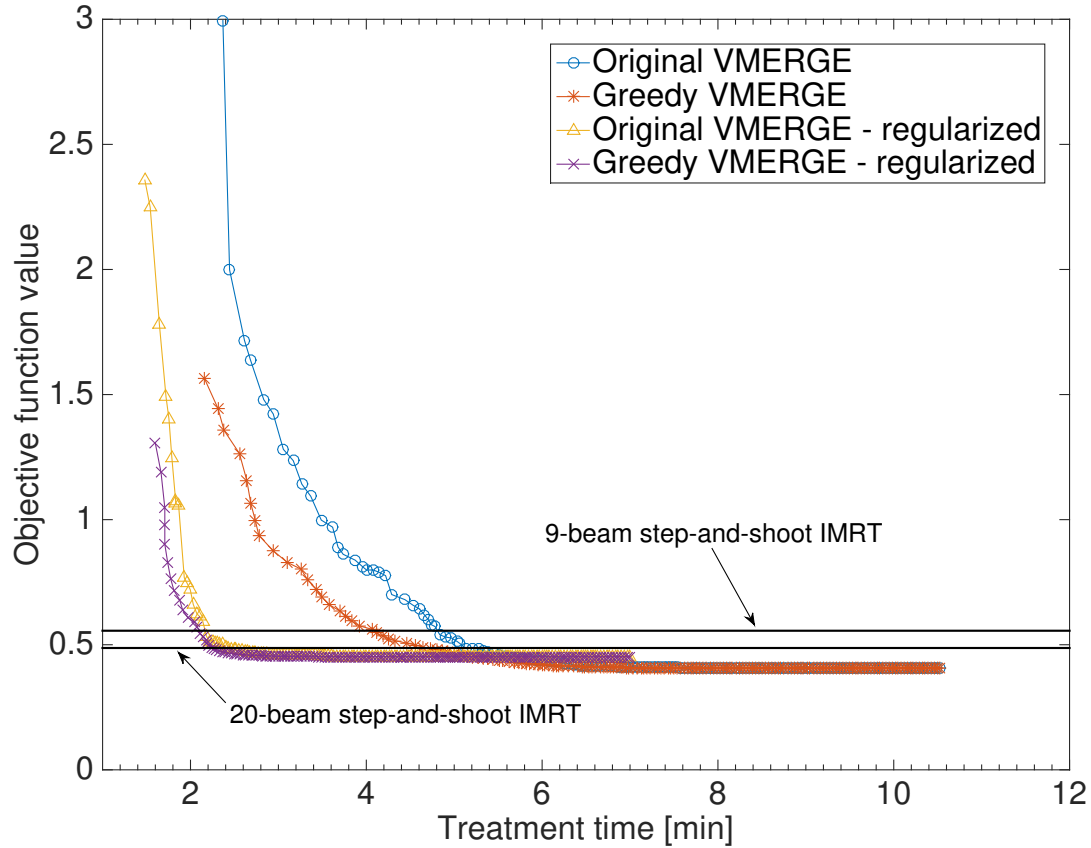


Figure 2.5: A comparison of treatment time versus plan quality trade-offs with the original and greedy VMERGE methods, with and without regularization, for the paraspinal case. The horizontal lines are objective function values for 9-beam and 20-beam step-and-shoot IMRT plans. As in the prostate case, using greedy merging improves the objective function value for the same treatment times, and adding SPG-regularization lowers the objective function value even more substantially. A treatment plan with the same plan quality as the optimal 20-beam IMRT plan can be delivered in just over two minutes.

objective function value than the original VMERGE method. In addition, the healthy organ sparing is comparable between the two VMERGE methods. Figure 2.5 provides a graphical summary of the results with the regularized solutions. Noting that it takes 360 seconds for the leaves to move across the 12 cm field 180 times, and allowing one minute for beam modulation, an upper bound of 420 seconds was used during the SPG-regularization step. As with the prostate case, the greedy variant and SPG-regularization provide a substantial reduction in treatment time for a given objective function value. The effect of the regularization is particularly notable. The plan quality of the 20-beam step-and-shoot IMRT plan is matched with a delivery time of just over two minutes.

2.5 Discussion

While the initial results are based on only two clinical cases, these results are both consistent and promising, indicating that the greedy merging strategy and especially the SPG-based regularization provide a substantial boost to the original VMERGE algorithm. The specific treatment times achieved are dependent on machine parameters such as leaf speed and patient geometry (especially the required field size), but the outcomes are qualitatively the same independent of these parameters. For example, we observed a similar reduction in treatment times (with the same plan quality) over the original VMERGE method assuming a 3 cm/s maximum collimator leaf speed.

There are a number of ways that the modified VMERGE algorithm could potentially be further improved. For example, the current optimization model explicitly assumes that the gantry speed is constant in each merged arc sector. A continuously variable gantry speed that can change within the merged arc sectors is more desirable, but it introduces another aspect to the optimization that adds considerable computational difficulty. Another limitation of the current approach is that its delivery time estimates rely on sliding window delivery, where all leaves have to move in the same direction across the entire field before they can simultaneously turn back. Allowing more general leaf trajectories would likely result in even lower treatment times [3, 24], but the nonconvexity of the underlying optimization models makes it particularly difficult to incorporate other types of leaf trajectories without increasing the plan optimization time to impractical levels.

Chapter 3

Optimizing nonuniformly fractionated treatments

In this chapter, we demonstrate that spatiotemporal fractionation schemes, in which the dose distribution can vary in each fraction, provide high-quality treatment plans and show an improvement in healthy tissue sparing over uniformly fractionated treatments. The optimization of spatiotemporal plans is based on a biological model, which introduces nonconvexity into the fluence map optimization. In five cases of liver cancer with distinct tumor geometries, the locally optimal spatiotemporal plans achieve substantial mean liver BED reduction over the optimal uniformly fractionated plans. The results indicate that spatiotemporal treatments can achieve substantial reductions in normal tissue dose and BED in the context of photon arc therapy.

3.1 Introduction

3.1.1 The linear-quadratic model for biologically effective dose

The effect of radiation on a tissue is given by the radiobiologic linear-quadratic (LQ) model for the number of cells that survive a radiation dose. In this model, there are two components to radiation-

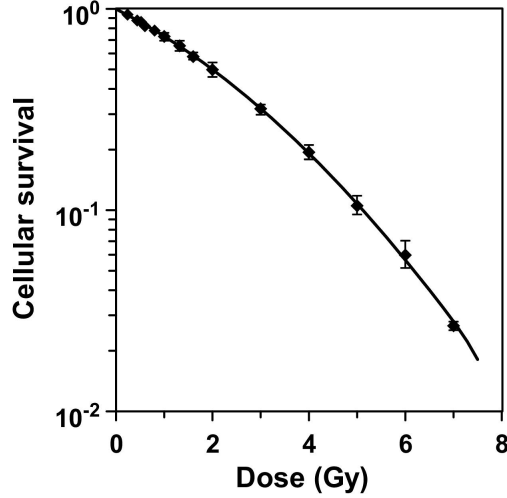


Figure 3.1: The cell survival curve given by the linear quadratic (LQ) model in (3.1). The solid line shows the estimation from the LQ model, and the data and error bars are from in-vitro experiments [18]. Reprinted with permission from Elsevier.

induced cell death, one that is proportional to the dose and one that is proportional to the dose squared. The fraction S of cells surviving dose d is given by

$$S = e^{(-\alpha d - \beta d^2)}, \quad (3.1)$$

where α and β are tissue-dependent parameters whose values are obtained from in-vitro experiments [32, 45]. A plot of the cell survival curves from the LQ model can be found in Figure 3.1. The effectiveness of a treatment is given by the logarithm of the surviving fraction of cells, and the equation is divided by α so the dimension is dose instead of dose squared [36]. This measure of effectiveness, called the *biologically effective dose* (BED), for a treatment with T fractions is given by

$$b = Td \left(1 + \frac{d}{\alpha/\beta} \right). \quad (3.2)$$

Recall from Chapter 1 that radiation treatments are fractionated by dividing the total dose into a series of smaller treatments. For a fixed total physical dose Td , the BED is minimized if the dose is split evenly into many fractions. This standard fractionation regimen is desired in normal tissues.

On the other hand, for a fixed total physical dose, the BED is maximized if all dose is delivered in a few fractions (called *hypofractionation*), which is desired in the tumor. Considering this inherent trade-off of fractionation decisions, it would seem ideal to simultaneously achieve hypofractionation in the tumor while splitting the dose to normal tissues evenly into many fractions. While this may appear unattainable at first glance, this goal can be approximately achieved using *spatiotemporal fractionation schemes*, in which the dose distribution in each fraction can vary. Spatiotemporal fractionation schemes deliver different dose distributions in different fractions in an attempt to minimize BED in healthy tissue and maximize BED in the tumor by hypofractionating parts of the tumor while delivering approximately identical doses to the surrounding tissue. Throughout this thesis, we use the terms *spatiotemporal fractionation* and *nonuniform fractionation* synonymously.

3.1.2 Relation to prior works

The clinical rationale for spatiotemporal fractionation to optimally exploit the dependence of cell survival on the fractionation scheme as described by the BED model was first proposed in the context of proton radiotherapy [97, 98, 110]. The potential benefit in IMPT comes from the fact that the dose in the entrance region of a proton beam is largely independent of the beam’s range, which provides some flexibility to modify the dose in the tumor without equally affecting the dose in the entrance region. (Recall the dose-depth curve for protons in Figure 1.3, which displays the Bragg peak and then the quick dose falloff after the peak.) Because of the differences in how proton doses and photon doses are absorbed, it was unclear until recently whether or not the benefit of spatiotemporal fractionation would translate to photon treatments. A proof-of-concept paper for spatiotemporal fractionation in photon therapy proposed that spatiotemporal fractionation may indeed provide a therapeutic advantage in arc therapy delivered with conventional photon beams [89]. The rationale for photon treatments is that distinct arc therapy plans can be created in such a way that each fraction delivers high single-fraction doses to complementary parts of the target volume while creating a similar dose bath in the surrounding normal tissue. This was demonstrated for the fractionated radiosurgery treatments of large arteriovenous malformations [92].

Prior research has addressed the problem of optimizing fractionation decisions, such as whether to hypofractionate or hyperfractionate, based on the BED model [50, 61, 94]. These works aim at maximizing the tumor BED subject to BED constraints to the normal tissue. It was shown that the optimal number of fractions depends on the α/β ratios of tumor and normal tissues, but also on the dose distribution. These works, however, all assume uniform fractionation, where the same dose distribution is delivered in all fractions.

The novelty in the idea of spatiotemporal fractionation lies in the fact that there is a potential advantage of delivering distinct dose distributions in different fractions, purely motivated by the cell response to fractionated treatments as described by the standard BED model. There are several extensions of the BED model that describe higher order biological effects such as incomplete repair of radiation damage between fractions [9, 103], repopulation of tumors over the course of treatment [16], the effect of chemotherapeutic agents [79], and reoxygenation of hypoxic tumors [109]. It was found that some of these models give rise to more complex fractionation schemes, i.e. varying doses per fraction. However, it is difficult to quantify the importance of any of these higher order effects in the clinic. Consequently, the role of such models to guide fractionation decisions in clinical practice has been limited. Instead, spatiotemporal fractionation as described in this thesis is purely based on the basic fractionation effect, whose existence and clinical relevance is undoubted.

Another approach in which different dose distributions in each fraction may arise is *adaptive radiotherapy* (ART) [108], sometimes called adaptive fractionation therapy [56] or dynamic radiotherapy [52, 53]. This technique uses feedback information, such as changes in the patient anatomy, to modify the treatment plan during the course of a fractionated treatment (e.g., replanning to compensate for tumor shrinkage in lung or head-and-neck cancer) [84, 106]. ART mostly deals with adapting treatment plans to compensate for geometric changes of the patient [51, 52]. Formally, different treatment plans are delivered in different fractions; however, ART in the context of geometric changes does not aim at purposefully delivering different tumor doses after replanning. To the contrary, it aims at reestablishing the paradigm of uniform fractionation.

3.2 A mathematical model for optimal spatiotemporal fractionation

3.2.1 Nonuniform fractionation using the BED model

The generalization of the BED model in (3.2) to fractionated treatments with T fractions is the *cumulative* BED defined as

$$b = \sum_{t=1}^T \left(d_t + \frac{d_t^2}{(\alpha/\beta)} \right), \quad (3.3)$$

where d_t is the physical dose delivered in fraction t , and α/β is a tissue-specific parameter; see, e.g., [66]. IMRT planning using BED can be performed analogously to conventional IMRT optimization, using similar objective functions and constraints in the treatment planning optimization model (1.1), but substituting BED in place of physical dose [97]. We define x_1, \dots, x_T to be the vectors of beamlet intensities delivered in fractions 1 through T , and V to be the set of voxels in the patient. Using the cumulative BED from (3.3), we obtain the spatiotemporal optimization problem:

$$\begin{aligned} \min_{x, d, b} \quad & F(b) \\ \text{s.t.} \quad & b_v = \sum_{t=1}^T \left(d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v} \right) \quad \forall v \in V \\ & Dx_t = d_t \quad t = 1, \dots, T \\ & x_t \geq 0 \quad t = 1, \dots, T, \end{aligned} \quad (3.4)$$

where D is the dose-influence matrix and the objective function F represents the desired clinical goals of target coverage, conformity, and organ sparing. We use a similar piecewise quadratic objective function that was first presented in Section 1.3.2, except the penalty is now a function of the BED b for a voxel or structure instead of the physical dose d . As before, the objective function F is a weighted sum of piecewise quadratic penalty functions that penalize BED above a prescribed threshold b_v^{hi} in voxel v , below a threshold b_v^{lo} , or above a prescribed mean BED m^{hi} in a given

structure. Let I^+ and I^- be the index sets of objectives that penalize BED above or below a prescribed threshold, and let I^m be the index set of objectives penalizing mean BED of a structure. Let V_i be the set of voxels involved in objective $i \in I$. We define the clinical goal F_i as the piecewise quadratic penalty function

$$F_i(b) = \begin{cases} \sum_{v \in V_i} (b_v - b_{iv}^{\text{hi}})_+^2 & \text{if } i \in I^+, \\ \sum_{v \in V_i} (b_{iv}^{\text{lo}} - b_v)_+^2 & \text{if } i \in I^-, \\ \left(\frac{1}{|V_i|} \sum_{v \in V_i} b_v - m_i^{\text{hi}} \right)_+^2 & \text{if } i \in I^m. \end{cases}$$

The overall objective function F is the weighted sum

$$F(b) = \sum_{i \in I} w_i F_i(b), \quad (3.5)$$

where I is the union of I^+ , I^- , and I^m , and each w_i reflects the relative importance of the i -th clinical goal.

There are typically protocols from common clinical practice that dictate the prescription values for dose to certain structures, but there are no such standards for BED prescriptions. Instead, we derive the prescription BED thresholds b_v^{hi} and b_v^{lo} from physical dose prescriptions. The prescription BED value is computed to be the BED that would result if the physical dose prescription were delivered in a uniformly fractionated treatment. For example, if the prescription value for the physical dose to a PTV structure (with $\alpha/\beta = 10$ Gy) is 50 Gy in 5 uniform fractions, the resulting BED of 100 Gy would be used as the BED prescription for the PTV in the spatiotemporal plan. We also employ the distance-dependent BED thresholds for voxels around the PTV as described in Section 1.3.2, using the same calculation from the physical dose thresholds.

Nonconvexity of nonuniform fractionation

Conventional fluence map optimization is a convex optimization problem when the clinical goals are modeled using a convex objective function F of the physical dose. In Section 3.2.2 we will show that despite the apparent nonconvexity introduced by the quadratic equality constraints, a BED-based optimization model for uniformly fractionated treatment plans is convex when the physical doses d are restricted to clinically relevant values and F is a piecewise quadratic penalty function similar to the dose-based objective functions. The nonuniform fractionation problem (3.4), however, is nonconvex. The nonconvexity of the formulation comes from the composition of the piecewise quadratic penalty function penalizing the underdose of the target and the function defining the BED. As a general rule, the composition of a convex nonincreasing function (such as our penalty function) with a convex quadratic function (such as the BED) is nonconvex, unlike the composition of a convex nondecreasing function (such as the functions used to penalize overdose) and the BED, which is always convex.

This nonconvexity is an inherent characteristic of the problem that cannot be eliminated by reformulating the model, as the following argument shows: by definition, permuting the fractions in an optimal treatment plan leads to another optimal plan, yet the average of these $T!$ treatment plans will be a plan with identical fractions, which in general will not be optimal. Thus, gradient-based optimization methods used to solve the uniform model can only yield locally optimal solutions for the nonuniform model. Later in this chapter, we will observe the nonconvexity of this problem in Figure 3.4, which shows that substantially different locally optimal spatiotemporal plans may exist for the same clinical case.

3.2.2 Uniform fractionation using the BED model

BED-based optimization can also be performed analogously for conventional, uniformly fractionated treatments. Throughout this work, we use a *uniform reference plan* as a benchmark to evaluate the benefit of nonuniform fractionation. We obtain the uniform reference plan by solving (3.4) with the additional constraint that $x_1 = \dots = x_T$. Eliminating the redundant variables, the uniform

reference plan is the optimal solution of the following problem:

$$\begin{aligned}
& \min_{x,d,b} F(b) \\
& \text{s.t.} \quad b_v = T d_v \left(1 + \frac{d_v}{(\alpha/\beta)_v} \right) \quad \forall v \in V \\
& \quad Dx = d \\
& \quad x \geq 0.
\end{aligned} \tag{3.6}$$

In a similar manner to [97], we verify that the BED-based optimization model for uniformly fractionated treatment plans is convex in the domain of “reasonable” doses d for all typical α/β values and number of fractions T . Because the BED is a monotone increasing function of the physical dose and the piecewise quadratic penalties are convex functions of the BED, the only potentially problematic term of the objective function is the underdose penalty of the form $f(b_v) = w(b_v^{\text{lo}} - b_v)_+^2 = w\left(b_v^{\text{lo}} - T\left(d_v + \frac{d_v^2}{\alpha/\beta}\right)\right)_+^2$ for a single voxel v . This is not a convex function of the physical dose. However, it can be seen that $\frac{\partial^2 f}{\partial d_v^2} > 0$ for physical dose values d_v that satisfy

$$d_v > -\frac{\alpha/\beta}{2} + \frac{1}{\sqrt{3}} \sqrt{\frac{(\alpha/\beta)b_v^{\text{lo}}}{T} + \frac{(\alpha/\beta)^2}{4}}. \tag{3.7}$$

For the parameter values used in the current work ($b_v^{\text{lo}} = 100$ Gy, $\alpha/\beta = 10$ for the tumor, and $T = 5$, as detailed later in Section 3.3.1), the above expression yields a lower bound of 3.66 Gy on the physical dose in order for the uniform model to be convex, which is not a binding constraint in the tumor, where the prescription BED b_v^{lo} and number of fractions T correspond to $d_v = 10$ Gy of physical dose per fraction. More generally, we may compare the lower bound in (3.7) with the prescribed per-fraction dose d_v^{lo} corresponding to a BED of b_v^{lo} delivered in T fractions, which is given by the formula

$$d_v^{\text{lo}} = -\frac{\alpha/\beta}{2} + \sqrt{\frac{(\alpha/\beta)b_v^{\text{lo}}}{T} + \frac{(\alpha/\beta)^2}{4}}. \tag{3.8}$$

This reveals that regardless of the parameters b_v^{lo} , α/β , and T , the underdose penalty is convex for per-fraction dose values satisfying

$$d_v > d_v^{\text{lo}}/\sqrt{3}, \quad (3.9)$$

which is reasonably required from every acceptable uniformly fractionated plan.

3.2.3 The constrained nonuniform model

While the objective function (3.5) can be useful to specify each of the clinical objectives and their relative priorities, it does not fully quantify the benefit of spatiotemporal fractionation. The optimal objective value has no physical interpretation, so if one treatment plan has a lower objective function value, we cannot specify in what way the treatment plan is “better” than another because an improvement in the objective function does not directly relate to a clinical outcome. Additionally, the benefit of spatiotemporal fractionation over conventional fractionation is distributed among the terms of the objective function, and the optimal solution does not realize the maximum benefit in any given objective.

In order to have a clinically interpretable objective function, we formulate an optimization problem in which we minimize only one of the clinical objectives of (3.5), say F_1 , while constraining the solution to be at least as good as the uniform reference plan with respect to each other objective F_i , ($i \in I, i \neq 1$). In the liver cases in Section 3.3, the primary goal is to lower the BED in the liver without compromising target coverage and conformity.

Let b^* be the BED delivered by the uniform reference plan, which is the globally optimal solution of (3.6), and let F_1 be the penalty function for nonzero mean BED in the liver. We modify (3.4) to include the constraints that the solution must be as good as the uniform reference plan with

respect to all objectives besides F_1 , and we obtain the constrained spatiotemporal fractionation problem:

$$\begin{aligned}
& \min_{x,d,b} && F_1(b) \\
& \text{s.t.} && F_i(b) \leq F_i(b^*) && i \in I, i \neq 1 \\
& && b_v = \sum_{t=1}^T (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) && \forall v \in V \\
& && Dx_t = d_t && t = 1, \dots, T \\
& && x_t \geq 0 && t = 1, \dots, T.
\end{aligned} \tag{3.10}$$

The first set of constraints ensures that the improvement in the objective F_1 is not at the cost of sacrificing the other clinical objectives; the computed spatiotemporal plan is either preferable or identical to the uniform reference plan with respect to every objective.

3.3 Application to liver tumors

We examined the benefit of spatiotemporal fractionation on two-dimensional slices of five clinical liver cases with distinct geometries. The first three cases feature centrally-located lesions within the liver; Case 1 has a large lesion, Case 2 has a small lesion, and Case 3 has two separate lesions within the liver. In each of these cases the liver is the main dose-limiting organ. In Case 4, the tumor abuts the chest wall, and Case 5 is a challenging geometry where both the chest wall and the bowel are dose-limiting and need to be included in the treatment plan optimization model. The patient geometries can be found in Figure 3.2.

3.3.1 Experimental setup

The BED prescription and α/β ratios

Five-fraction treatments were planned for all cases. To derive the upper and lower BED thresholds b^{hi} and b^{lo} , typical prescription doses and normal tissue constraints for 5-fraction liver treatments

were converted into BED values using the clinically accepted values of $\alpha/\beta = 10$ Gy in the tumor and $\alpha/\beta = 4$ Gy in all normal tissues [32,67]. The prescription lower bounds b^{lo} for the GTV and PTV were chosen to be 100 Gy and 72 Gy, respectively, which correspond to 50 Gy and 40 Gy of physical dose delivered in 5 fractions assuming uniform fractionation. We also included an upper bound b^{hi} of 115.5 Gy BED in the GTV and 100 Gy BED in the PTV, corresponding to 55 Gy and 50 Gy of physical dose delivered in 5 uniform fractions.

We minimize mean liver BED ($m^{\text{hi}} = 0$) to spare the healthy liver tissue. In addition, to achieve a conformal dose distribution, we implement a distance-dependent quadratic penalty function for voxels outside the PTV. We impose a linear falloff of BED from 175 Gy at the edge of the PTV to a BED of 15 Gy in 3 cm. This corresponds to a physical dose falloff from 50 Gy to 10 Gy in 5 equal fractions. In Cases 4 and 5, with the chest wall near the tumor, we highly penalize dose exceeding $b^{\text{hi}} = 96.2$ Gy BED (35 Gy of physical dose in 5 fractions) to the chest wall. Additionally, in Case 5, where the GI tract is also dose-limiting, we add an upper bound of 75 Gy BED (30 Gy of physical dose in 5 fractions) to the affected sections of the GI tract.

Computing the uniformly fractionated reference plan

The BED-based fluence map optimization problem (3.6) was solved to obtain optimal beamlet weights for a uniformly fractionated reference plan with 21 equispaced beams, which approximates a high-quality VMAT plan by delivering the modulated beams in small arc sectors [68]. The number of arc sectors could be reduced by using techniques such as the merging of adjacent fluence maps described in Chapter 2. Dose-influence matrices were calculated for the 21 beams using the dose calculation algorithm implemented in CERR version 5.2 [31]. In this work, we limited the computations to two-dimensional slices of the patient voxels and optimized beamlet weights for a single row of the beamlet grid. As discussed in Section 3.2.2, although the uniform plan optimization problem is defined using a combination of convex and nonconvex constraints, when the beamlet weights x and physical doses d are restricted to clinically relevant values, the problem is convex, and the globally optimal solution can be computed by local optimization algorithms.

L-BFGS-B is a quasi-Newton method to solve bound constrained optimization problems, which are problems that aim to minimize a differentiable nonlinear function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ subject to the constraints $\ell \leq x \leq u$. L-BFGS-B is a descendant of the BFGS algorithm, which minimizes an unconstrained objective function f with iterates of the form

$$x_{k+1} = x_k - \alpha_k H_k \nabla f_k,$$

where α_k is an appropriately-chosen step size and H_k is an approximation of the inverse Hessian of f that is updated in every iteration using the gradients from previous iterations. When the number of variables is large, the cost to store the approximate Hessian matrix may be problematic. The limited-memory variant L-BFGS stores a set of vectors that contain curvature information from the past m iterations (where m is a user-defined parameter), and the Hessian approximation is defined implicitly instead of storing the entire matrix. With the addition of the bound constraints $\ell \leq x \leq u$, the L-BFGS-B algorithm is based on the gradient projection method and uses the limited-memory BFGS Hessian approximation [112]. In our experiments, the optimization was performed with Matlab using the L-BFGS-B solver [111], and the runtimes to obtain the globally optimal uniformly fractionated solutions were less than 10 seconds on a standard desktop computer for all of the cases.

Computing the spatiotemporally fractionated plans

After computing the uniform reference plan, we computed nonuniformly fractionated treatment plans in which the beamlet weights and corresponding dose distributions are not the same in each of the five fractions. As described in Section 3.2.3, we solve (3.10) to minimize the mean BED in the liver, subject to the constraints that the solution must be at least as good as the uniform reference plan with respect to all other objectives. In this work, we used the interior point algorithm implemented in Matlab's built-in optimizer `fmincon` [58] to compute locally optimal solutions of these nonconvex problems.

Interior point methods, which were originally proposed for linear programming [49], find solutions to constrained nonlinear programs using a barrier function that approaches infinity as a point approaches the boundary of the feasible region. For a general nonlinearly constrained optimization problem

$$\begin{aligned} \min \quad & f(x) \\ \text{s.t.} \quad & h(x) = 0 \\ & g(x) \leq 0, \end{aligned}$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $h : \mathbb{R}^n \rightarrow \mathbb{R}^t$, and $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are smooth functions, the barrier problem is

$$\begin{aligned} \min \quad & f(x) - \mu \sum_{i=1}^m \ln s_i \\ \text{s.t.} \quad & h(x) = 0 \\ & g(x) + s = 0, \end{aligned} \tag{3.11}$$

for some value of $\mu > 0$ [19]. Path-following interior point methods iterate through the interior of the feasible region with a sequence of points x_μ , which converge to an optimal solution of the constrained optimization problem as μ approaches zero. For a fixed value of μ , an iterate x_μ is computed by solving a system of optimality conditions for the barrier problem (3.11) with Newton's method [35].

In the current work, we initialize the optimization with a perturbation of the fluence maps of the uniform reference plan. The beamlet weights were multiplied by independent random factors drawn uniformly from the interval $[0, 2]$, and the quality of the computed local optimal solutions did not appear to be sensitive to the specific way the plans were perturbed. Runtimes for computing a locally optimal solution of (3.10) ranged from 15 minutes to 4 hours.

3.3.2 Results

Figure 3.3 shows a comparison of the dose distributions of the uniform reference plan with locally optimal nonuniformly fractionated solutions for Case 1. In Figure 3.4, several additional locally optimal solutions for Case 1 are presented to demonstrate that the locally optimal solutions par-

Table 3.1: Summary of mean liver BED reductions from spatiotemporal fractionation and sparing factors for each of the five cases. The sparing factor is a value that determines the dependence of the optimal fractionation schedule of a fixed dose distribution on the patient geometry. The remarkable benefit seen in Case 2 agrees with the fact that the sparing factor is substantially lower in this case than in the other cases.

	Case description	Mean liver BED [Gy]			Sparing factor
		conventional	spatiotemporal	reduction	
1	Large central lesion	84.54	75.87	12.8%	0.6663
2	Small lesion	26.14	19.47	34.3%	0.3830
3	Two small lesions	59.54	50.24	18.5%	0.5879
4	Lesion abutting ch. wall	47.51	38.65	22.9%	0.5289
5	Lesion abutting GI tract	88.67	77.38	14.6%	0.7028

tition the tumor differently but maintain similar overall treatment quality. Similar figures for the remaining cases can be found in Appendix A.

Table 3.1 summarizes the computed bounds and benefits observed in all five cases. In each case, the mean liver BED in the nonuniform plans is substantially lower than in the uniform plans. Case 1 displays the smallest reduction in the mean liver BED with an approximately 13% improvement, and Case 2 has the largest reduction of approximately 34%. Note that these improvements are achieved without sacrificing any other clinical objective, as by definition, the spatiotemporal plans computed by solving (3.10) are at least as good as the uniform reference plan with respect to every objective.

The effectiveness of a spatiotemporal treatment can be quantified by its equivalent dose DEQ5, defined via

$$\text{BED} = \sum_{t=1}^5 \left[d_t + \frac{d_t^2}{\alpha/\beta} \right] = \text{DEQ5} + \frac{(\text{DEQ5})^2/5}{\alpha/\beta}$$

which yields

$$\text{DEQ5} = 5 \left[-\frac{\alpha/\beta}{2} + \sqrt{\left(\frac{\alpha/\beta}{2} \right)^2 + \frac{\alpha/\beta}{5} \sum_{t=1}^5 \left[d_t + \frac{d_t^2}{\alpha/\beta} \right]} \right].$$

The DEQ5 is the physical dose distribution that achieves the same BED as the spatiotemporal treatment, if the treatment were to be delivered in five uniform fractions. Using DEQ5 has the

advantage that the spatiotemporal plan can be directly compared to the physical dose distribution of the uniform reference plan. The comparison is shown in Figure 3.3e and in the remaining figures of Appendix A. In the visualization of the difference between the uniform and DEQ5 dose distributions, we observe that the nonuniformly fractionated plans maintain the same tumor BED as the uniformly fractionated plans while reducing the mean dose and BED in healthy liver tissue. This reduction can also be seen by comparing the DVH curves of the physical dose and DEQ5 of the two plans. For Case 1, these curves are included in Figure 3.5. The DVH curves show that there is substantial dose reduction in the liver (both the DEQ5 and physical dose curves shift to the left), and that this is a consequence of lower physical dose (but nearly identical DEQ5) delivered to the tumor.

The optimal number of fractions

In order to fairly evaluate the benefit of spatiotemporal fractionation, we must consider the potential gain from changing the fractionation schedule in addition to altering the physical dose distribution. We recomputed the optimal uniformly fractionated treatment plans with up to five fractions by solving the uniform model (3.6) for $T = 1, \dots, 5$. In all of the cases, the five-fraction uniform plans had the lowest mean liver BED among all computed uniform plans. Hence, Table 3.1 compares the 5-fraction spatiotemporal plans against the best uniform plans with up to five fractions.

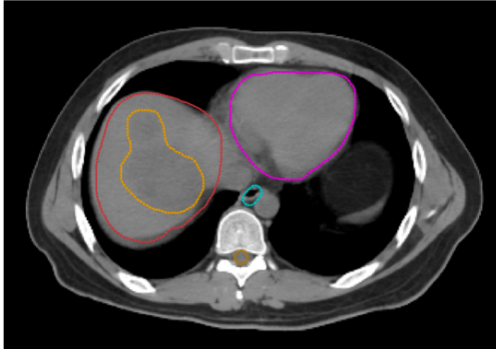
This agrees with previous work on the dependence of optimal fractionation schedules on the patient geometry and dose distribution. When the goal is to minimize the mean dose to a single dose-limiting parallel organ such as the liver, the optimal fractionation schedule is a function of the *effective sparing factor* $\bar{\delta}$ which was introduced independently by [94] and [50]. Using the notation $(\alpha/\beta)_N$ and $(\alpha/\beta)_T$ for the α/β ratios for the normal tissue and tumor, respectively, it was shown that if $\bar{\delta} > \frac{(\alpha/\beta)_N}{(\alpha/\beta)_T} = 0.4$, then increasing the number of fractions is optimal, and if $\bar{\delta} < 0.4$, then lowering the number of fractions is optimal. In Case 2, the effective sparing factor is approximately 0.4, while in the remaining four cases, the sparing factor is well above 0.4. This provides an explanation for the observation that the benefit of spatiotemporal fractionation was largest for Case 2; the treatment quality of a uniformly fractionated 5-fraction treatment and a single-fraction treatment

is approximately equal. Therefore, achieving a benefit through spatiotemporal fractionation relies to a lesser extent on achieving near-uniform fractionation in the normal liver.

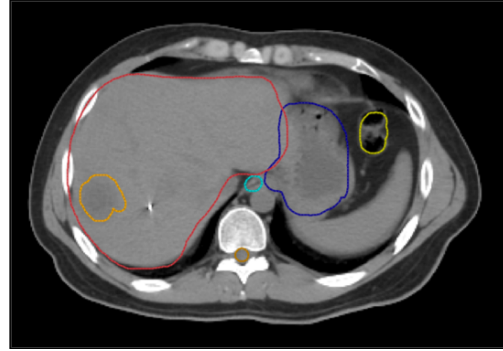
3.4 Discussion

The results indicate that spatiotemporal treatments can achieve substantial reductions in both BED and physical dose to the liver and other normal tissue. The approximately 13-35% mean liver BED reduction is consistent with the benefit of spatiotemporal fractionation observed in previous work for IMPT and large arteriovenous malformations [92].

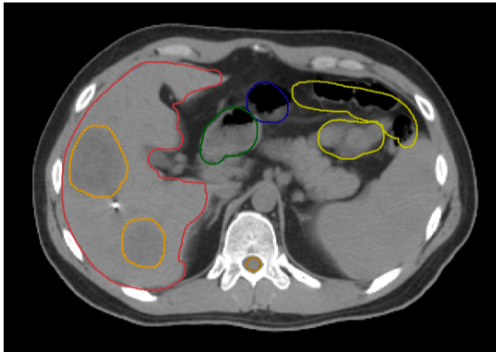
As mentioned in Section 3.3, the treatment plans in this work have a higher number of beams than a conventional IMRT plan to approximate a VMAT plan, where we expect nonuniform fractionation to display the most benefit. The nonuniformly fractionated plans that are most effective in lowering BED and total physical dose are those in which the tumor dose is hypofractionated while the dose to surrounding tissue is fractionated. In other words, these plans deliver a high single-fraction dose to parts of the tumor during each fraction and a consistent lower dose to the liver and other healthy tissue throughout all of the fractions. VMAT treatments are particularly suitable for delivering such a nonuniformly fractionated plan because the gantry rotation causes a lower dose to be spread over the entire arc, which helps preserve the healthy tissue.



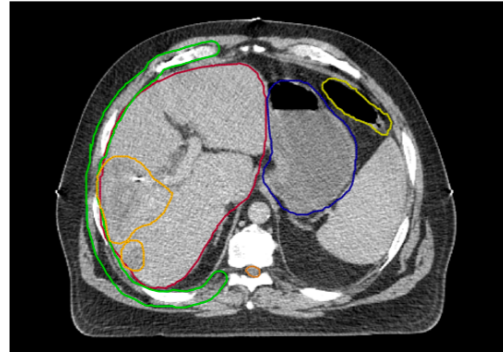
Case 1: Large central lesion



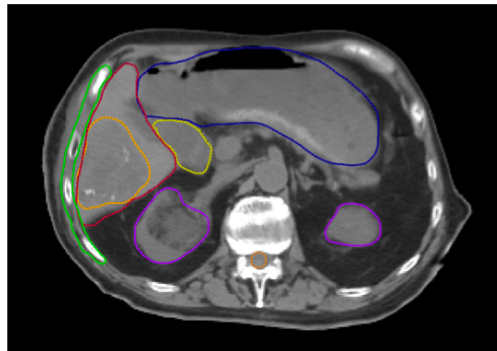
Case 2: Small lesion



Case 3: Two small lesions



Case 4: Lesion abutting chest wall



Case 5: Lesion abutting GI tract

Figure 3.2: Patient geometries for each of the five cases. The structures contoured on the CT scan are the liver (red), GTV (orange), chest wall (light green), large and small bowel (yellow), stomach (dark blue), esophagus (light blue), kidneys (purple), spinal cord (light brown), heart (magenta), and duodenum (dark green).

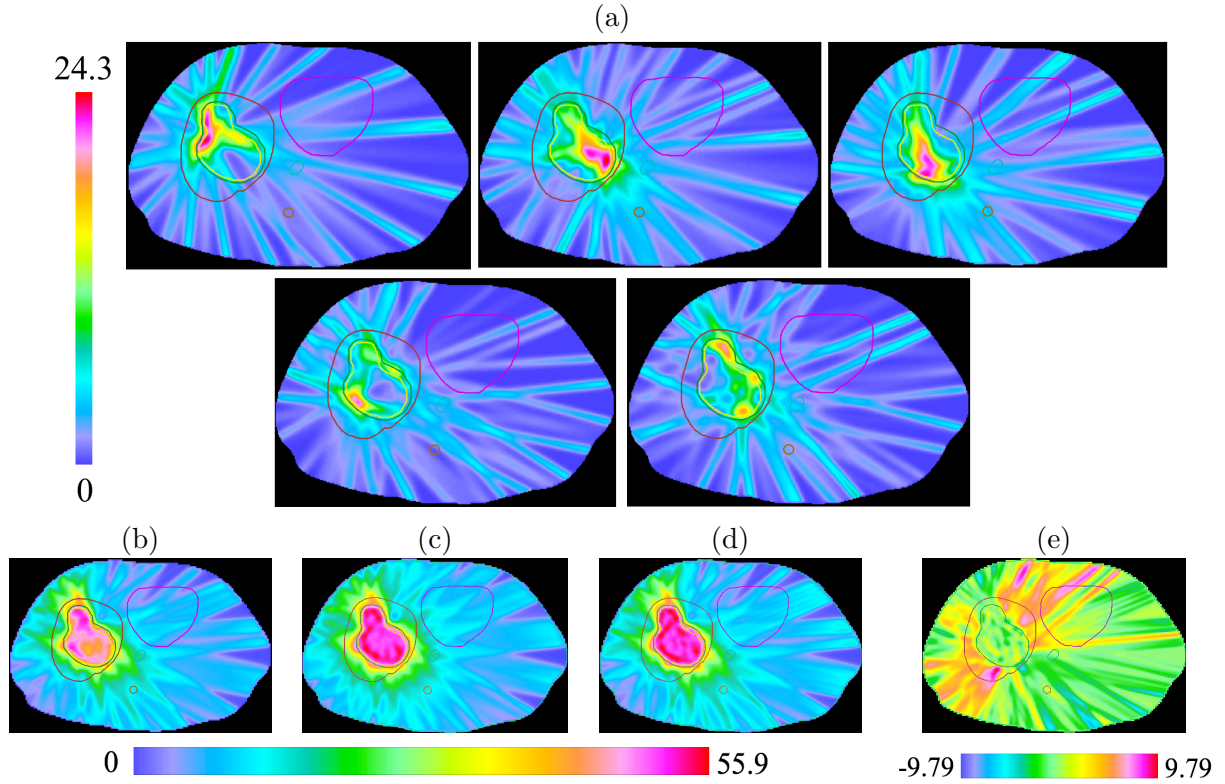


Figure 3.3: Dose distributions for Case 1, a large central lesion within the liver. (a) Physical dose distributions in each of the five fractions show that the nonuniformly fractionated treatment hypofractionates different parts of the tumor. (b) Total physical dose delivered throughout the nonuniformly fractionated treatment. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan, which is the uniform plan that is isoeffective in delivering the same BED as the nonuniformly fractionated plan. (e) The difference between the physical dose in the uniform plan and the DEQ5 for the nonuniform plan, or (c) minus (d). This shows that the spatiotemporal plans reduce dose in the healthy liver and in the entrance region of the those beams that expose the liver the most. All numerical quantities shown are in [Gy].

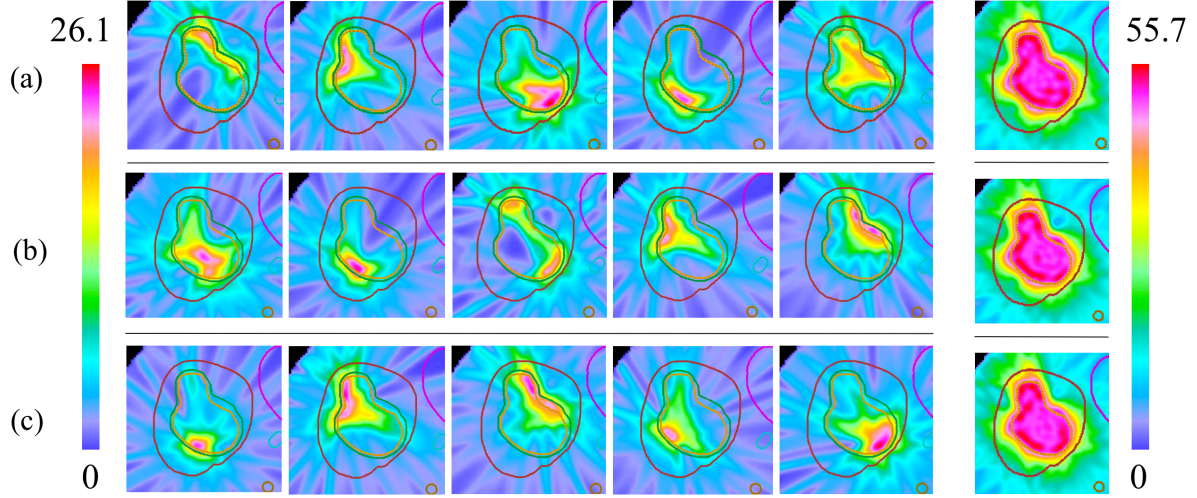
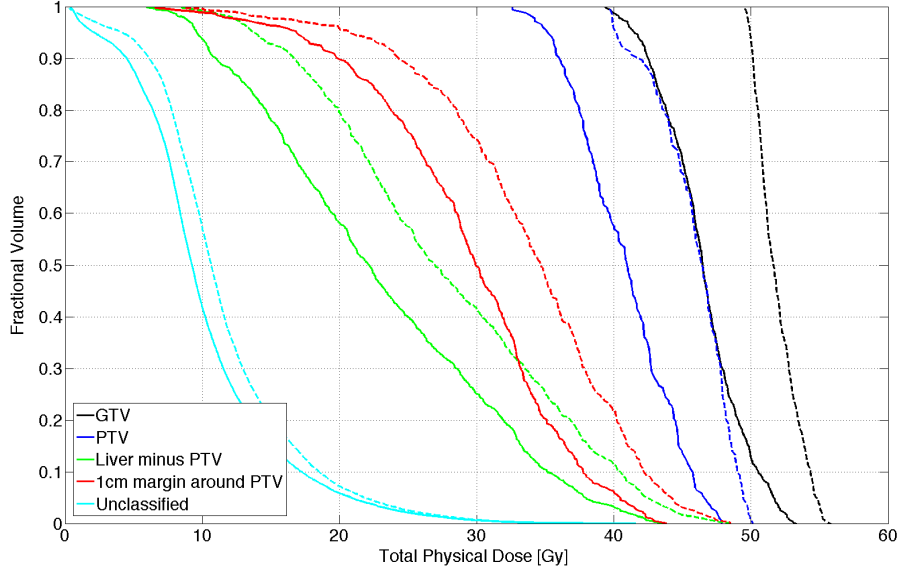
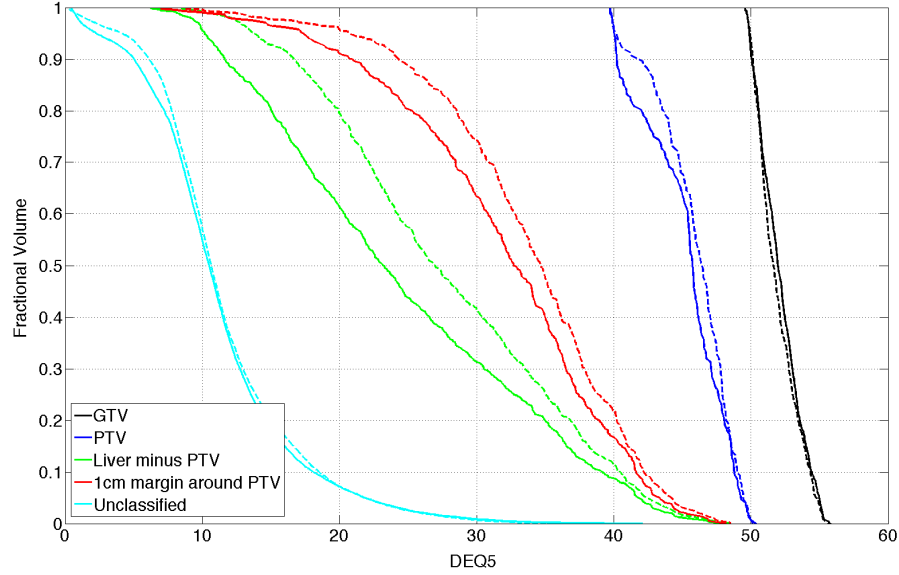


Figure 3.4: Rows (a), (b), and (c) are physical dose distributions from three additional locally optimal 5-fraction nonuniform treatments for Case 1, which is the same clinical liver case displayed in Figure 3.3. They are all locally optimal solutions of the model (3.4). The first five panels of each row are the dose distributions in the five nonuniform fractions, and the last panel on each row is the equivalent dose DEQ5. The solutions exhibit the same pattern: different subregions of the tumor receive a high single-fraction dose in different fractions. Note that the emergent “partitions” in the optimized plans are a result of the optimization. This pattern supports the rationale that the benefit of spatiotemporal fractionation is a result of hypofractionating parts of the tumor while maintaining a consistent low dose in the surrounding tissue. The difference in the hypofractionated regions in each solution also demonstrates that several qualitatively different locally optimal spatiotemporal treatments may exist for the same case. All numerical quantities shown are in [Gy].



(a)



(b)

Figure 3.5: Dose-volume histogram for (a) total physical dose and (b) DEQ5 for various structures in Case 1. The dashed lines are the values from the uniformly fractionated reference plan and the solid lines are from the spatiotemporal plan. These curves indicate that the spatiotemporal plan achieves an overall reduction in physical dose, as observed by the lines shifted to the left in (a). Additionally, in (b) we note that the spatiotemporal plan maintains DEQ5 in the tumor while reducing dose to the healthy liver tissue. To assess the conformity of the dose around the PTV, we included a curve for the voxels in a 1-cm margin around the PTV.

Chapter 4

Bounding the achievable benefit of spatiotemporal fractionation

In the previous chapter, we demonstrated that treatment plans optimized with a nonuniformly fractionated treatment scheme can result in a lower BED in healthy tissue while maintaining the tumor coverage of a conventional uniformly fractionated plan. However, the spatiotemporal treatment plans presented in the last chapter were only locally optimal solutions of nonconvex optimization problems, and we had little information about the optimality of these solutions on the global scale. In this chapter, we begin by summarizing the theory of how to use semidefinite programming to formulate a convex relaxation of a nonconvex quadratically constrained quadratic program (QCQP). Next, we present a semidefinite programming relaxation of the constrained spatiotemporal fractionation model. We present the results of solving the semidefinite programming relaxation with the five clinical liver cases used in the last chapter, which allows us to bound the maximum achievable benefit of spatiotemporal fractionation. We conclude by discussing some limitations of this methodology and potential future directions of research.

4.1 Semidefinite programming relaxations of general QCQPs

A *semidefinite program* (SDP) is an optimization problem of the form

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & c^\top x \\ \text{s.t.} \quad & A(x) \succcurlyeq 0, \end{aligned}$$

where $c \in \mathbb{R}^n$ and $A : \mathbb{R}^n \rightarrow \mathbb{S}^n$ is an affine mapping [101]. (The notation \mathbb{S}^n represents the cone of $n \times n$ symmetric matrices.) The expression $A(x) \succcurlyeq 0$ is called a *linear matrix inequality*, which imposes that $A(x)$ must be positive semidefinite. Similarly, the notation $A \succ 0$ is used to indicate the matrix A is positive definite. A formulation with a linear objective function and a single linear matrix inequality may seem restrictive, but it can be shown that many convex constraints can be equivalently written as linear matrix inequalities [6, Ch. 3]. Hence, many convex optimization problems can be reformulated into equivalent semidefinite programs.

The idea of formulating a convex relaxation of a general QCQP using semidefinite matrices was first proposed by Shor, using Lagrangian duality to obtain a lower bound on the optimal value [82]. Semidefinite relaxations were then applied to a special case of combinatorial optimization problems using the technique of lifting the problem to write the quadratic inequalities as linear inequalities in a higher-dimensional space [55]. Later, Vandenberghe and Boyd used similar techniques to develop a SDP relaxation for general QCQP that is the dual of Shor's relaxation; these two relaxations yield the same lower bound on the optimal value of a general QCQP [101]. In the remainder of this section, we will describe a procedure for formulating a SDP relaxation of a general QCQP. Consider the following general QCQP with inequality constraints:

$$\begin{aligned} \min_x \quad & x^\top A_0 x + b_0^\top x + c_0 \\ \text{s.t.} \quad & x^\top A_i x + b_i^\top x + c_i \leq 0 \quad i = 1, \dots, m, \end{aligned} \tag{4.1}$$

where $A_i \in \mathbb{S}^n$ (but not necessarily positive semidefinite), $b_i \in \mathbb{R}^n$, and $c_i \in \mathbb{R}$ for all $i = 0, 1, \dots, m$. Recalling that $x^\top Ax = \text{tr}(Axx^\top)$ and introducing an auxiliary matrix variable $X = xx^\top$, we can write the QCQP as

$$\begin{aligned} \min_{x, X} \quad & \langle A_0, X \rangle + b_0^\top x + c_0 \\ \text{s.t.} \quad & \langle A_i, X \rangle + b_i^\top x + c_i \leq 0 \quad i = 1, \dots, m \\ & X = xx^\top, \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ represents the matrix inner product $\langle A, X \rangle = \text{tr}(A^\top X)$. The final constraint, $X = xx^\top$ is equivalent to imposing that $X \succcurlyeq 0$ and $\text{rank}(X) = 1$. The constraint $X \succcurlyeq 0$ is a convex constraint, but the constraint $\text{rank}(X) = 1$ is not. We relax the constraint $X - xx^\top = 0$ into $X - xx^\top \succcurlyeq 0$; this makes the feasible region larger, so the solution to the relaxation is a lower bound for the solution of the original QCQP.

We now have an optimization problem with m linear constraints and one quadratic matrix inequality:

$$\begin{aligned} \min_{x, X} \quad & \langle A_0, X \rangle + b_0^\top x + c_0 \\ \text{s.t.} \quad & \langle A_i, X \rangle + b_i^\top x + c_i \leq 0 \quad i = 1, \dots, m \\ & X - xx^\top \succcurlyeq 0. \end{aligned}$$

To reformulate this as a semidefinite program, we use the Schur complement to write the quadratic matrix inequality as a linear matrix inequality. For a block matrix $X = \begin{pmatrix} A & B \\ B^\top & C \end{pmatrix}$ with an invertible block A , the Schur complement of X is $C - B^\top A^{-1} B$. If we assume that $A \succ 0$, then we have

$$X \succcurlyeq 0 \iff C - B^\top A^{-1} B \succcurlyeq 0$$

[17]. Using this fact, we can rewrite the last constraint as the linear matrix inequality $\begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix} \succcurlyeq 0$. Then a relaxation of the QCQP (4.1) is the semidefinite program

$$\begin{aligned} \min_{x, X} \quad & \langle A_0, X \rangle + b_0^\top x + c_0 \\ \text{s.t.} \quad & \langle A_i, X \rangle + b_i^\top x + c_i \leq 0 \quad i = 1, \dots, m \\ & \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix} \succcurlyeq 0. \end{aligned} \tag{4.2}$$

The optimal objective function value of (4.2) is a lower bound for the solution of (4.1).

4.2 A mathematical model for bounding the maximum achievable benefit of spatiotemporal fractionation

We now apply the concepts from the previous section to formulate a convex optimization model for bounding the maximum achievable benefit from nonuniform fractionation. We begin by writing problem (3.4)–(3.5) as a quadratic optimization problem with quadratic constraints.

Recall that the objective function for the optimization of spatiotemporally fractionated plans is a weighted sum of three types of penalty functions (see Section 3.2.1). We introduce the auxiliary variables p_{iv} for the overdose penalties $(b_v - b_{iv}^{\text{hi}})_+$, the auxiliary variables q_{iv} for the underdose penalties $(b_{iv}^{\text{lo}} - b_v)_+$, and the auxiliary variables r_i for the mean dose penalties $\left(\frac{1}{|V_i|} \sum_{v \in V_i} b_v - m_i^{\text{hi}}\right)_+$. Note that in general, $(y)_+$ is the smallest number z satisfying the inequalities $z \geq y$ and $z \geq 0$. We impose similar inequality constraints on the auxiliary variables p_{iv} , q_{iv} , and r_i :

$$\begin{aligned} p_{iv} &\geq b_v - b_{iv}^{\text{hi}} & \forall i \in I^+, \forall v \in V_i \\ p_{iv} &\geq 0 \\ q_{iv} &\geq b_{iv}^{\text{lo}} - b_v & \forall i \in I^-, \forall v \in V_i \\ q_{iv} &\geq 0 \end{aligned}$$

$$r_i \geq \frac{1}{|V_i|} \sum_{v \in V_i} b_v - m_i^{\text{hi}} \quad \forall i \in I^{m+}$$

$$r_i \geq 0.$$

When these auxiliary variables are minimized, they are equal to the values of the positive part functions.

We eliminate the variables b_v from the problem using the equality

$$b_v = \sum_{t=1}^T (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}),$$

and we arrive at the following quadratically constrained quadratic programming (QCQP) formulation of the nonuniform fractionation problem:

$$\begin{aligned}
\min_{x, d, p, q, r} \quad & \sum_{i \in I^+} \sum_{v \in V_i} w_i p_{iv}^2 + \sum_{i \in I^-} \sum_{v \in V_i} w_i q_{iv}^2 + \sum_{i \in I^m} w_i r_i^2 \\
\text{s.t.} \quad & p_{iv} \geq -b_{iv}^{\text{hi}} + \sum_{t=1}^T (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) \quad \forall i \in I^+, \forall v \in V_i \\
& p_{iv} \geq 0 \quad \forall i \in I^+, \forall v \in V_i \\
& q_{iv} \geq b_{iv}^{\text{lo}} - \sum_{t=1}^T (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) \quad \forall i \in I^-, \forall v \in V_i \\
& q_{iv} \geq 0 \quad \forall i \in I^-, \forall v \in V_i \\
& r_i \geq \frac{1}{|V_i|} \sum_{v \in V_i} \left(\sum_{t=1}^T d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v} \right) - m_i^{\text{hi}} \quad \forall i \in I^m \\
& r_i \geq 0 \quad \forall i \in I^m \\
& Dx_t = d_t \quad t = 1, \dots, T \\
& x_t \geq 0 \quad t = 1, \dots, T.
\end{aligned} \tag{4.3}$$

Because the physical dose d_{vt} is a linear function of the beamlet intensities x_t , the d_{vt} variables can be eliminated from the formulation, and all the inequality constraints can be seen as quadratic

inequalities in the primary decision variables x_1, \dots, x_T . By introducing the auxiliary matrix variables $X_t = x_t x_t^\top$, the quadratic inequalities in (4.3) can be written as linear inequalities in x_t and X_t . For instance, on the right-hand side of the first set of inequalities we substitute

$$-b_{iv}^{\text{hi}} + \sum_{t=1}^T (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) = \sum_{t=1}^T \left\langle \begin{pmatrix} 1 & x_t^\top \\ x_t & X_t \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{\text{hi}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & \frac{1}{(\alpha/\beta)_v} D^\top e_v e_v^\top D \end{pmatrix} \right\rangle,$$

where e_v is the characteristic vector with $e_v = 1$ in the v -th position and zeros elsewhere. (That is, $e_v^\top D$ is simply the v -th row of the D matrix.) The resulting optimization model has a convex quadratic objective function and only linear constraints, aside from the nonconvex quadratic equations $X_t = x_t x_t^\top$ for $t = 1, \dots, T$. As described in the previous section, we obtain a convex relaxation of this problem by replacing this set of equations by the weaker convex constraint $X_t - x_t x_t^\top \succcurlyeq 0$ and, using the Schur complement, this constraint is equivalent to the convex linear matrix inequality $\begin{pmatrix} 1 & x_t^\top \\ x_t & X_t \end{pmatrix} \succcurlyeq 0$.

Since in the original nonconvex model we have $X_t = x_t x_t^\top$ and $x_t \geq 0$, it is clear that each X_t is component-wise nonnegative. Hence, we can add the inequalities $X_t \geq 0$ to the convex relaxation. (This component-wise inequality should not be confused with the linear matrix inequality $X_t \succcurlyeq 0$.) Note that even though the inequality $X_t \geq 0$ is redundant in the original model, $X_t \geq 0$ is not redundant in the convex relaxation and tightens the bound.

A further simplification is possible. Since the relaxation is convex and symmetric in the fractions, we can assume without loss of generality that at the optimum we have $x_1 = \dots = x_T$ and $X_1 = \dots = X_T$. Thus, we can eliminate the variables corresponding to the different fractions and use only a single variable x and X in place of each x_t and X_t . (This shows that our convex relaxation does not distinguish between the uniformly and nonuniformly fractionated models, although the bound does depend on the number of fractions.)

Finally, using the shorthand

$$C_v = \frac{1}{(\alpha/\beta)_v} D^\top e_v e_v^\top D,$$

we arrive at the following convex relaxation of (4.3):

$$\begin{aligned}
& \min_{x, X, p, q, r} && \sum_{i \in I^+} \sum_{v \in V_i} w_i p_{iv}^2 + \sum_{i \in I^-} \sum_{v \in V_i} w_i q_{iv}^2 + \sum_{i \in I^m} w_i r_i^2 \\
& \text{s.t.} && p_{iv} \geq T \left\langle \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{\text{hi}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^+, \forall v \in V_i \\
& && p_{iv} \geq 0 && \forall i \in I^+, \forall v \in V_i \\
& && q_{iv} \geq -T \left\langle \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{\text{lo}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^-, \forall v \in V_i \\
& && q_{iv} \geq 0 && \forall i \in I^-, \forall v \in V_i \\
& && r_i \geq \frac{T}{|V_i|} \sum_{v \in V_i} \left\langle \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{m_i^{\text{hi}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^m \\
& && r_i \geq 0 && \forall i \in I^m \\
& && \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix} \succcurlyeq 0, \quad x \geq 0, \quad X \geq 0.
\end{aligned} \tag{4.4}$$

The optimal objective function value of this problem is a lower bound for the minimum value of the spatiotemporal fractionation problem (3.4)-(3.5).

Note that to obtain (4.4), we linearized every quadratic constraint in (4.3), even though all of them were convex except for those involving q_{iv} . The bound from (4.4) is tighter than what could be obtained by simply replacing the concave quadratic constraints with a convex relaxation and keeping the convex quadratic constraints intact. This is because unlike the linearizations of the concave quadratics (which are relaxations), the linearizations of the convex quadratics are tighter than the original constraints. Our derivation of (4.4) shows that the convex model as a whole is indeed a relaxation of (4.3) despite the fact that this cannot be seen when we compare the models constraint by constraint.

4.2.1 Bounding the maximum benefit in a given objective

In Section 3.2.3 we formulated a constrained version of the spatiotemporal fractionation model to prioritize a single clinical objective, which allowed us to maximize the benefit of the prioritized objective. The convex relaxation of the constrained nonuniform model (3.10) can be derived analogously to how the relaxation (4.4) of the nonuniform model (3.4) was obtained. Consistent with the application to liver tumors, we assume that the mean dose objective

$$F_1(b) = \left(\frac{1}{|V_1|} \sum_{v \in V_1} b_v - m_1^{\text{hi}} \right)_+^2$$

is the primary objective, where V_1 is the set of healthy liver voxels. We minimize the auxiliary variable r_1 for the mean liver BED, and, similar to the constraints $F_i(b) \leq F_i(b^*)$, $\forall i \in I$, $i \neq 1$ in (3.10), we impose the constraints on the auxiliary variables p_{iv} and q_{iv} . Then the convex optimization model bounding the minimum value of F_1 from below is the following:

$$\begin{aligned}
& \min_{x, X, p, q, r} && r_1 \\
& \text{s.t.} && \sum_{v \in V_i} p_{iv}^2 \leq F_i(b^*) && \forall i \in I^+ \\
& && \sum_{v \in V_i} q_{iv}^2 \leq F_i(b^*) && \forall i \in I^- \\
& && r_i \leq F_i(b^*) && \forall i \in I^m, i \neq 1 \\
& && p_{iv} \geq T \left\langle \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{\text{hi}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^+, \forall v \in V_i \\
& && p_{iv} \geq 0 && \forall i \in I^+, \forall v \in V_i \\
& && q_{iv} \geq -T \left\langle \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{\text{lo}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^-, \forall v \in V_i
\end{aligned} \tag{4.5}$$

$$\begin{aligned}
q_{iv} &\geq 0 & \forall i \in I^-, \forall v \in V_i \\
r_i &\geq \frac{T}{|V_i|} \sum_{v \in V_i} \left\langle \begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{m_i^{\text{hi}}}{T} & \frac{e_v^\top D}{2} \\ \frac{D^\top e_v}{2} & C_v \end{pmatrix} \right\rangle & \forall i \in I^m \\
r_i &\geq 0 & \forall i \in I^m \\
\begin{pmatrix} 1 & x^\top \\ x & X \end{pmatrix} &\succcurlyeq 0, \quad x \geq 0, \quad X \geq 0.
\end{aligned}$$

In the next section, we describe how the first two convex quadratic constraints can be reformulated into an equivalent linear matrix inequality [101], so without loss of generality we can include them in the semidefinite program.

4.2.2 Solution methods

In our implementation of the convex relaxations, we reformulate the models to include both semidefinite constraints and second-order cone constraints. The *second-order cone* in \mathbb{R}^{K+1} (also known as the *Lorentz cone* or the *quadratic cone*) is defined as

$$\mathcal{Q}_{K+1} = \{(x_0, \bar{x}) \in \mathbb{R} \times \mathbb{R}^K \mid x_0 \geq \|\bar{x}\|\},$$

and a second-order cone constraint requires a vector to be in the second-order cone. Second-order cone constraints are a special case of semidefinite constraints because it can be shown that any second-order cone constraint can be written as a semidefinite constraint. In particular, $(x_0, \bar{x}) \in \mathcal{Q}_{K+1}$ if and only if the arrow-shaped matrix

$$\text{Arw}(x_0, \bar{x}) = \begin{pmatrix} x_0 & x_1 & \cdots & x_K \\ x_1 & x_0 & & \\ \vdots & & \ddots & \\ x_K & & & x_0 \end{pmatrix}$$

(where all entries not on the diagonal, first row, or first column are zero) is positive semidefinite [1]. Semidefinite programming solvers often have better performance with second-order cone constraints exploiting the compact structure of vectors in the second-order cone rather than using the larger matrix representation. In our implementation, instead of writing the first quadratic constraint in the model (4.5) as a linear matrix inequality, we represent the constraint $\sum_{v \in V_i} p_{iv}^2 \leq F_i(b^*)$ as the second order cone constraint

$$(\sqrt{F_i(b^*)}, p_{i,1}, \dots, p_{i,|V_i|}) \in \mathcal{Q}_{|V_i|+1}.$$

Interior point methods, which are used to solve both linear and nonlinear programs, can also be generalized to solve semidefinite programming problems. Interior point methods are currently considered the state of the art for solving SDPs; they can solve SDPs in polynomial time and perform well for many applications [6, Ch. 4–5]. Interior point methods for SDPs are implemented in widely available optimization software such as Sedumi [85] and Mosek [63]. In our experiments we used Mosek to solve the semidefinite programs (4.4) and (4.5).

4.3 Results for clinical liver cases

To evaluate the quality of the locally optimal spatiotemporal solutions, we solved the SDP relaxation for each of the five liver cases from Section 3.3 in the previous chapter. See Section 3.3.1 for details about the cases and the prescriptions. The lower bounds obtained from solving the SDP relaxations can be found in Table 4.1; the first few columns in Table 4.1 are replicated from Table 3.1 to serve as a reference for the lower bound.

Bounds on the maximum achievable liver BED reduction

The lower bounds obtained from the SDP relaxation are close to the mean liver BED values achieved by the locally optimal nonuniform solutions. In Case 1, spatiotemporal fractionation reduced the mean liver BED from 84.5 Gy in the uniform reference plan to 75.9 Gy, which is a reduction of 8.7

Table 4.1: Summary of mean liver BED reductions from spatiotemporal fractionation and lower bounds for mean liver BED. The “gap closed” values provide a measure of how close the local optimal solutions are to achieving the lower bound on the mean liver BED; see Eq. (4.6) for the definition.

	Case description	Mean liver BED [Gy]			Gap closed
		conventional	spatiotemporal	reduction	
1	Large central lesion	84.54	75.87	12.8%	77.7%
2	Small lesion	26.14	19.47	34.3%	88.2%
3	Two small lesions	59.54	50.24	18.5%	80.8%
4	Lesion abutting ch. wall	47.51	38.65	22.9%	89.9%
5	Lesion abutting GI tract	88.67	77.38	14.6%	96.9%

Gy from the uniform plan. The lower bound from the SDP relaxation for Case 1 is 73.4 Gy, which means that no treatment plan can achieve a reduction of more than 11.2 Gy from the uniform reference plan. It is important to note that the value of 73.4 Gy is only a lower bound, and we are not able to guarantee the existence of a nonuniform solution that achieves this value. The mean liver BED value for the true globally optimal plan may be anywhere between 75.9 and 73.4 Gy.

We compared the mean liver BED reduction in the locally optimal nonuniform plan with the bound on the maximum possible reduction by computing the fraction of the gap between the mean liver BED in the uniform reference plan and the lower bound from the SDP relaxation that the nonuniform plan closes. Let m_{ref} and m_{sp} denote the mean liver BED in the uniform reference plan and the spatiotemporal plan, respectively, and notate the lower bound provided from the semidefinite programming relaxation by m_{SDP} . Then the *gap closed* is the ratio

$$\frac{m_{\text{ref}} - m_{\text{sp}}}{m_{\text{ref}} - m_{\text{SDP}}}. \quad (4.6)$$

This gap represents the percentage of the maximum possible benefit that the spatiotemporal plans achieve. If the gap closed is near 100%, then the spatiotemporal plan is very near the SDP lower bound; similarly, if the gap closed is small, then there is still a considerable distance between the mean liver BED in the spatiotemporal plan and the lowest possible mean liver BED. For example,

in Case 1, the nonuniform plan's reduction of 8.7 Gy is 78% of the upper bound on the maximum possible reduction of 11.2 Gy. In the five clinical liver cases, the nonuniformly fractionated plans closed 78-97% of the gap between the mean liver BED of the uniform reference plan and the bound on the lowest achievable mean liver BED.

4.4 Discussion

The results in this chapter are limited to the optimization of treatment plans for two-dimensional slices of clinical liver cases, and the main challenge in extending this study to three-dimensional cases is the computational complexity of solving the SDP relaxations in Section 4.2. In the nonuniform fractionation optimization problem there is at least one variable for each beamlet in the fluence maps and at least one constraint for every voxel associated with a quadratic penalty function. For reference, the largest optimization problem out of the five two-dimensional slices in this work contains 282 beamlets in a single fraction, while a three-dimensional clinical liver case may have over 3,000 beamlets. In the SDP relaxation, with the introduction of the matrix variables X , the number of variables is roughly squared. Semidefinite programs of this size cannot be solved for three-dimensional cases in a reasonable amount of time with the available off-the-shelf software. Even with the two-dimensional case, the solver took up to 7 hours to solve the SDP. It is an interesting direction for future research to devise customized algorithms for the solution of these large-scale semidefinite optimization problems that arise for the three-dimensional cases. For now, the results obtained for the two-dimensional slices suggest that the locally optimal solutions computed for the nonuniform fractionation problem are close to global optimality.

Chapter 5

Robust spatiotemporal fractionation schemes in the presence of patient setup uncertainty

In this chapter, we address the concern of random patient setup uncertainty in the context of spatiotemporal fractionation. A stochastic optimization model is used to optimize spatiotemporally fractionated plans using expected penalties for deviations from prescription values, and a discrete probability distribution is defined to characterize the random setup error occurring in each fraction. For the optimization of uniform plans, the expected penalties are computed exactly by exploiting the symmetry of the fractions, and for the spatiotemporal plans, quasi-Monte Carlo sampling is used to approximate the expectation. Using five clinical liver cases it is demonstrated that spatiotemporally fractionated treatment plans maintain the same robust tumor coverage as a stochastic uniform reference plan and exhibit a reduction in the expected mean BED of the uninvolved liver.

5.1 Introduction

In previous chapters, we demonstrated the remarkable benefit of spatiotemporal fractionation over conventional uniform fractionation. The benefit from spatiotemporally fractionated photon plans is a result of their characteristic pattern of delivering high single-fraction doses to small parts of the tumor. The resulting sharp dose gradients need to be meticulously aligned in different fractions in order to avoid potentially compromising tumor control. As the regions receiving high dose in each fraction are not predetermined, but are the result of the treatment plan optimization, this issue cannot be addressed using margins around the treated volumes. In this chapter, we study the impact of random setup uncertainty on spatiotemporal treatments. We utilize stochastic optimization techniques to directly incorporate setup uncertainty in the treatment planning optimization problem. When this uncertainty is properly accounted for, spatiotemporal photon plans can be computed that achieve substantial normal tissue sparing compared to conventional plans while being as robust against random setup uncertainty as their conventionally fractionated counterparts.

The problem of patient setup uncertainty is not new in the context of IMRT and IMPT. For a recent, comprehensive review of the subject, the reader is referred to [90]. Previous work utilizes stochastic optimization techniques to handle uncertainties in radiotherapy treatments by constraining or optimizing the expected values of quantities such as quadratic penalty functions [10, 96], expected values of objective functions based on tumor control probability [105], or objective functions based on an approximation of the expected cumulative dose [34]. Another stochastic optimization approach is to use a mean and variance constraint for each voxel to ensure robust target coverage [21]. In the current work we use a scenario-based model with expectations of penalty functions to account for the uncertainty in patient positioning, and use, when necessary, scenario sampling to approximate the expected values. An important difference between spatiotemporal and conventional fractionation lies in the importance of random errors, i.e. setup errors that are different from fraction to fraction. In conventional fractionation, random errors have a minor impact compared to systematic errors, i.e. a setup error that is the same in each fraction. Since spatiotemporal fraction-

ation delivers a distinct dose distribution in each fraction, random errors are both more important and computationally more challenging to account for.

5.2 Treatment plan optimization for spatiotemporal fractionation under uncertainty

5.2.1 Modeling dose delivery uncertainty

Random setup errors, along with other sources of uncertainty in dose delivery, can be modeled as uncertainty in the dose-influence matrix D . We use the notation Ω for the set of possible setup errors, or *scenarios*, and D^ω for the realization of D in scenario $\omega \in \Omega$. Under random setup uncertainty, the patient positioning error is potentially different in each fraction, and it is convenient to think of each scenario as a T -dimensional vector $\omega = (\omega_1, \dots, \omega_T)$. The associated probability function is denoted by P . For simplicity, in this chapter we will work with finitely many discrete scenarios, that is, $\Omega \subset \mathbb{R}^T$ is a finite set, and each scenario ω has a positive probability $P(\omega)$. Our choice of distribution P for the patient setup uncertainty is described in more detail later in Section 5.3.1.

Recall the cumulative BED equation given in (3.3), which is based on the linear-quadratic model of a tissue's response to radiation over the course of multiple fractions. Now, with the dose-influence matrix being a random variable, the physical dose and the cumulative BED in each voxel also become random variables. The cumulative BED for voxel v in scenario ω is

$$b_v^\omega = \sum_{t=1}^T \left(d_{vt}^{\omega_t} + \frac{(d_{vt}^{\omega_t})^2}{(\alpha/\beta)_v} \right). \quad (5.1)$$

As a result, the plan quality that we aim to optimize is also random. Following [95] and many others, we quantify plan quality by the expected value of an appropriately chosen penalty function that penalizes deviation from prescribed BED values, such as the piecewise quadratic penalty functions in Section 3.2.1. For example, to ensure that the cumulative BED delivered to a particular voxel remains above a prescribed value b_v^{lo} with sufficiently high probability, we may minimize (or

constrain) the quantity

$$\mathbb{E}_P \left[(b_v^{\text{lo}} - b_v)_+^2 \right]. \quad (5.2)$$

To lighten the notation, for the rest of the chapter, we shall drop the subscript P from the expected values, as the probability distribution is always the same.

5.2.2 Stochastic optimization of uniformly fractionated treatments

To ensure a fair comparison between spatiotemporal fractionation and uniform fractionation, uniformly fractionated treatment plans are optimized using the same BED-based prescriptions and the same stochastic model of random setup uncertainty as the spatiotemporal plans. The benefit of spatiotemporal fractionation can be evaluated by comparing the spatiotemporal plans to their robust uniformly fractionated counterparts. In our optimization models, all clinical objectives are implemented using standard piecewise quadratic penalty functions. Structures that are sufficiently distant from the target volumes are likely to be nearly or entirely unaffected by random setup uncertainty; it is sufficient to consider only the nominal scenario in the objectives involving these structures. The objectives concerning target structures and nearby volumes are implemented using expected piecewise quadratic penalties.

Let I^+ and I^- denote the index sets of the clinical objectives associated with structures that are assigned an expected penalty for the BED exceeding the prescription amount b^{hi} or BED falling below a prescription amount b^{lo} , respectively. Let I^m be the index set for the clinical objectives penalizing mean BED exceeding m^{hi} . Similarly, let \bar{I}^+ , \bar{I}^- , and \bar{I}^m denote the index sets of the clinical objectives that involve nominal penalty values. Let the set I be the union of the six aforementioned index sets. Let V_i be the set of voxels associated with the objective $i \in I$, with $V = \cup_{i \in I} V_i$ denoting the set of all patient voxels considered in the optimization model. Lastly, let $\bar{\omega}$ denote the nominal scenario.

Piecewise quadratic penalty functions F_i are used to mathematically define the clinical objectives, either for the nominal scenario $\bar{\omega}$ or the expected value over all scenarios:

$$F_i(b) = \begin{cases} \sum_{v \in V_i} (b_v^{\bar{\omega}} - b_{iv}^{\text{hi}})_+^2 & \forall i \in \bar{I}^+ \\ \sum_{v \in V_i} (b_{iv}^{\text{lo}} - b_v^{\bar{\omega}})_+^2 & \forall i \in \bar{I}^- \\ \left(\frac{1}{|V_i|} \sum_{v \in V_i} b_v^{\bar{\omega}} - m_i^{\text{hi}} \right)_+^2 & \forall i \in \bar{I}^m \\ \mathbb{E} \left[\sum_{v \in V_i} (b_v - b_{iv}^{\text{hi}})_+^2 \right] & \forall i \in I^+ \\ \mathbb{E} \left[\sum_{v \in V_i} (b_{iv}^{\text{lo}} - b_v)_+^2 \right] & \forall i \in I^- \\ \mathbb{E} \left[\left(\frac{1}{|V_i|} \sum_{v \in V_i} b_v - m_i^{\text{hi}} \right)_+^2 \right] & \forall i \in I^m. \end{cases} \quad (5.3)$$

Using this notation, the uniformly fractionated reference plans are obtained by solving optimization problems of the form

$$\begin{aligned} \min_{x, d, b} \quad & \sum_{i \in I} q_i F_i(b) \\ \text{s.t.} \quad & b_v^\omega = \sum_{t=1}^T \left(d_{vt}^{\omega_t} + \frac{(d_{vt}^{\omega_t})^2}{(\alpha/\beta)_v} \right) \quad \forall v \in V, \quad \omega = (\omega_1, \dots, \omega_T) \in \Omega \\ & d_t^\omega = D^{\omega_t} x \quad \forall \omega \in \Omega, \quad t = 1, \dots, T \\ & x \geq 0, \end{aligned} \quad (5.4)$$

with the penalty weights $q_i \geq 0$ reflecting the relative importance of each clinical objective. The variables d_{vt}^ω and b_v^ω can be eliminated using the equality constraints; this yields an optimization problem with the fluence x as the only variable, and the only constraint is the nonnegativity of the fluence.

In the stochastic model, the expected value of the underdose penalty is the probability-weighted linear combination of deterministic quadratic underdose penalties; thus, a sufficient condition for the convexity of the stochastic model (5.4) is that the per-fraction dose values satisfy (3.9) in every scenario.

5.2.3 Stochastic optimization of spatiotemporal treatments

It is straightforward to formulate an optimization model analogous to (5.4) to compute an optimal robust spatiotemporal plan, with only two small changes to the model (5.4). First, the fluence map in each fraction t needs to be a separate vector of decision variables x_t , and then the corresponding random physical dose d_t^ω is computed as $d_t^\omega = D^{\omega_t} x_t$. One way to compare spatiotemporal treatment plans with uniform reference plans would be to optimize both plans for the same objective function (i.e. the same set of objective weights q_i). However, improvements in terms of quadratic penalty functions are difficult to interpret. In the application to liver tumors in Section 5.4, we instead quantify the benefit of spatiotemporal fractionation via the mean BED reduction in the healthy liver for a fixed BED delivered to the tumor. To that end, we minimize the objective function corresponding to the primary clinical objective (the mean liver BED) while constraining the remaining penalty function values to be at least as good as the penalty values in the uniform reference plan.

Next, we provide the mathematical details for this formulation. Without loss of generality, let V_1 denote the voxel set associated with the primary clinical objective and let F_1 be the corresponding penalty function. Let b^* denote the BED distribution of the uniformly fractionated reference plan, with $F_i(b^*)$ the value of the i -th penalty function evaluated with the uniform plan. For all clinical objectives aside from the primary clinical objective, we constrain the spatiotemporally fractionated treatment plan to have a smaller or equal penalty function value than the uniform reference plan.

Then the optimization model for spatiotemporally fractionated plans can be written as follows:

$$\begin{aligned}
& \min_{x,d,b} F_1(b) \\
& \text{s.t.} \quad F_i(b) \leq F_i(b^*) \quad \forall i \in I, i \neq 1 \\
& \quad b_v^\omega = \sum_{t=1}^T \left(d_{vt}^{\omega_t} + \frac{(d_{vt}^{\omega_t})^2}{(\alpha/\beta)_v} \right) \quad \forall v \in V, \quad \omega = (\omega_1, \dots, \omega_T) \in \Omega \\
& \quad d_t^\omega = D^{\omega_t} x_t \quad \forall \omega \in \Omega, \quad t = 1, \dots, T \\
& \quad x_t \geq 0 \quad t = 1, \dots, T.
\end{aligned} \tag{5.5}$$

From the same reasoning outlined in Section 3.2.1 for the deterministic model, this stochastic spatiotemporal optimization model is nonconvex, and the solutions to this model can only be certified to be locally optimal.

5.3 Modeling random patient setup uncertainty

We model uncertainty in patient positioning by considering scenarios in which the patient is shifted slightly from the nominal position.

5.3.1 Probability distributions of patient position

First, let us consider a one-dimensional probability distribution of the patient's position when the setup uncertainty is restricted to one axis of motion. For computational convenience, we assume that the patient is shifted an integer number of voxels on either side of the nominal position. Let p_s be the probability that the patient's position is s voxels away from the nominal position. (Thus, p_0 is the probability of the patient being in the nominal position.) We assume that the probabilities p_s satisfy

$$p_s = \gamma^{|s|} p_0 \quad s \in S \tag{5.6}$$

for some parameter $\gamma \in [0, 1]$ and a finite set $S \subset \mathbb{Z}$ symmetric about zero. In particular, the patient positioning error is assumed to be symmetric and bounded, with its mode in the nominal

position, and the probability of the patient being in the nominal position is given by

$$p_0 = \frac{1}{\sum_{s \in S} \gamma^{|s|}}.$$

The parameter γ allows us to consider a spectrum of probability distributions. As γ increases, the probabilities of larger errors increase. When $\gamma = 0$, the probability of the nominal scenario is 1, and all other shifts have probability zero. At the other extreme, $\gamma = 1$, the nominal scenario and all shifts have an equal probability of occurring. In all of our experiments, $S \subseteq \{-2, -1, 0, 1, 2\}$.

We assume that the patient positioning errors along each axis are independent. Considering $|S|$ scenarios for the error along each of the n axes, the number of scenarios for the patient position in one fraction is $|S|^n$. In a treatment with T fractions, the number of scenarios is $|S|^{n \times T}$. For example, in some of our computational experiments, a five-fraction treatment plan is computed that allows the patient position to be shifted up to two voxels from the nominal position in either direction, in two dimensions. The number of distinct scenarios in these experiments is $5^{2 \times 5} \approx 10^7$.

5.3.2 Computing the expected penalty values

Because we use a finitely supported probability distribution, the expected value of the penalty function F can be computed exactly as a finite sum over all scenarios:

$$\mathbb{E}[F(b_v)] = \sum_{\omega \in \Omega} P(\omega) F(b_v^\omega), \quad (5.7)$$

where $P(\omega)$ is the probability of scenario ω . However, the large number of scenarios in the spatiotemporal models becomes prohibitive during optimization, as the time required to compute the penalty function values and their gradients is proportional to the number of scenarios. Later in Section 5.4.1, we will discuss the static dose cloud assumption that is adopted throughout this chapter. This assumption allows us to accelerate the function evaluations for the expected penalties because the dose in scenario ω is a permutation of the dose in the nominal scenario, and the

expected penalty functions can be computed efficiently using permutations of the nominal dose calculation. The calculation of the gradients of the penalty functions is the bottleneck in the computations because they require multiplication by the permuted dose-influence matrices instead of simply permuting the components of the dose vector.

Approximating penalty values with quasi-Monte Carlo sampling

The time required to compute the function values and gradients of (5.7) at least once during every iteration of an optimization algorithm can be mitigated by taking a sample $\tilde{\Omega} \subseteq \Omega$ containing $|\tilde{\Omega}|$ scenarios and approximating the expectation with

$$Q_N = \frac{1}{|\tilde{\Omega}|} \sum_{\omega \in \tilde{\Omega}} F(b_v^\omega). \quad (5.8)$$

A simple method to obtain the sample $\tilde{\Omega}$ is Monte Carlo sampling, which draws a set of N elements according to the probability distribution and approximates the expectation with the estimator above. The approximation error of a Monte Carlo sample is given by

$$\left| \mathbb{E}[F(b_v)] - Q_N \right|, \quad (5.9)$$

which converges to zero at a rate of $\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$. One advantage of Monte Carlo sampling is that this bound on the approximation error is dimension-free, which makes it a good choice for problems with high-dimensional uncertainty.

Quasi-Monte Carlo (QMC) methods utilize more evenly distributed sequences of points to approximate the expectation, and generally the approximation error of QMC sequences has better convergence than that of Monte Carlo approximations when the dimension of the uncertainty is sufficiently small. In our application, the dimension of the uncertainty is nT (with $n = 2$ and $T = 5$ for our experiments), so we use QMC integration, rather than Monte Carlo, to sample the scenarios and approximate the expected penalty values. In particular, we employ a rank-one lattice integration rule. An N -point rank-one lattice is a point set in $[0, 1]^s$, where s is the dimension of

the uncertainty (with $s = nT$ in the current work), of the form

$$P_N = \left\{ \left(\frac{zi}{N} \right) \bmod 1, i = 0, \dots, N-1 \right\}, \quad (5.10)$$

where $z \in \mathbb{R}^s$ is called the *generating vector* [54]. Rank one lattices are low discrepancy point sets whose approximation error converges at a rate of $\mathcal{O}(N^{-1}(\log N)^s)$ [64]. A generating vector for a rank-one lattice can be constructed using a greedy component-by-component algorithm to minimize the worst-case approximation error [83].

An advantage of QMC lattices is that they can be randomized using the Cranley-Patterson procedure to obtain statistical bounds on the approximation error [30]. This procedure adds a random shift Δ that is uniformly distributed in $[0, 1)^s$ to each vector in the lattice. With this randomization, the lattice P_N becomes

$$\tilde{P}_N(\Delta) = \left\{ \left(\frac{zi}{N} + \Delta \right) \bmod 1, i = 0, \dots, N-1 \right\}. \quad (5.11)$$

The estimator (5.8) with a randomized rank-one lattice sample is unbiased because F evaluated on each random lattice point is an unbiased estimator of the true expected value, and the average of these unbiased estimators is also unbiased [54].

Computing exact penalty values in uniformly fractionated plans

We note that for uniformly fractionated plans, scenario sampling is not necessary to evaluate the expected penalty values. In a uniformly fractionated treatment, every permutation of the T shifts in a scenario $\omega = (\omega_1, \dots, \omega_T)$ will yield the same cumulative BED because the same fluence map is delivered in each fraction, and the cumulative BED is not dependent on the order of the fractions. Thus, after a penalty is evaluated for a particular sequence of T shifts, its weight can be adjusted to account for all permutations of that shift sequence. Formally, let $\hat{\Omega} \subset \Omega$ be the set of scenarios $\hat{\omega} = (\hat{\omega}_1, \dots, \hat{\omega}_T)$ such that every scenario $\omega = (\omega_1, \dots, \omega_T) \in \Omega$ is a permutation of exactly one scenario in $\hat{\Omega}$; furthermore, for each $\hat{\omega} \in \hat{\Omega}$, let $N(\hat{\omega})$ be the number of scenarios in Ω that are

a permutation of $\hat{\omega}$, and let $P(\omega)$ be the probability of scenario ω . Then the expectation of the penalty function F satisfies

$$\mathbb{E}[F(b_v)] = \sum_{\omega \in \Omega} P(\omega) F(b_v^\omega) = \sum_{\hat{\omega} \in \hat{\Omega}} N(\hat{\omega}) P(\hat{\omega}) F(b_v^{\hat{\omega}}). \quad (5.12)$$

Even with T as little as 5, $\hat{\Omega}$ is substantially smaller than Ω , which vastly reduces the number of function evaluations required to compute the exact expected value. For example, in our experiments with 2-dimensional patients and patient positioning errors up to 2 voxels in either direction along both axes, the number of scenarios can be reduced from $5^{10} \approx 10^7$ to 118 755. This eliminates the need for scenario sampling in the optimization models for uniformly fractionated treatments because the expected penalty value over all scenarios can be exactly computed in each iteration of the optimization algorithm.

5.4 Numerical experiments with clinical liver cases

We computed uniformly fractionated and spatiotemporally fractionated treatment plans for the same five two-dimensional slices of liver tumors from the previous two chapters. In this chapter, we considered patient setup errors in the two-dimensional transverse plane; we did not consider superior-inferior motion of the patient that would move the voxels out of the transverse slice.

5.4.1 Clinical liver cases and prescriptions

In each of the five cases, an α/β ratio of 10 was used for the target structure and an α/β ratio of 4 was used for all healthy tissues. Five-fraction treatments were optimized to be consistent with common clinical practice. The optimization models incorporated the following clinical objectives:

- Minimize the expected mean BED in the uninvolved liver, defined as the liver minus the GTV.

- Penalize the expected shortfall from 100 Gy BED in the GTV. This corresponds to 50 Gy of physical dose delivered in 5 fractions.
- Minimize the nominal BED in the unclassified tissue.
- For Cases 4 & 5: Penalize excess of 96.25 Gy BED in the chest wall in the nominal scenario. This is equivalent to 35 Gy of physical dose in 5 fractions.
- For Cases 4 & 5: A conformal plan was desired, with a prescribed linear falloff from 175 Gy BED to 15 Gy BED in 3 cm in the nominal scenario. This is equivalent to a falloff from 50 Gy to 10 Gy of physical dose in 3 cm.
- For Case 5: Minimize the nominal BED in the kidneys.

BED excess and shortfall were penalized quadratically (from a threshold of 0 Gy if not specified otherwise), as detailed in Section 5.2.2.

The nominal dose-influence matrix D was obtained using the Quadrant Infinite Beam (QIB) dose calculation algorithm implemented in CERR version 5.2 [31]. Recall that the patient setup error is modeled as error in the dose-influence matrix, so each scenario $\omega \in \Omega$ has its corresponding matrix D^ω . Instead of repeating the dose calculation for every possible position of the patient to explicitly find each D^ω , we adopt the static dose cloud assumption, which is commonly used in treatment planning optimization [81, 90]. Under this assumption, the position of the patient does not affect the dose distribution, and if a shift of a patient occurs, the nominal dose distribution is delivered to the shifted patient. Each D^ω can be assembled from the rows of the nominal D with no additional dose calculation.

5.4.2 Probability distributions of patient shifts

In this work we considered a spectrum of probability distributions on three different supports. First, we considered probability distributions supported on 1- and 2-voxel shifts on a dose grid that was downsampled from the CT resolution by a factor of 2. The voxel sizes for these experiments

were $2.54 \text{ mm} \times 2.54 \text{ mm}$ for the first four cases and $2.18 \text{ mm} \times 2.18 \text{ mm}$ for Case 5. With other treatment sites in mind that have smaller setup uncertainties, we also ran a set of experiments without downsampling and using the full CT resolution with probability distributions supported on 1-voxel shifts. These experiments used a voxel size of $1.27 \text{ mm} \times 1.27 \text{ mm}$ for Cases 1-4 and $1.09 \text{ mm} \times 1.09 \text{ mm}$ for Case 5. In all of the experiments, we assumed that the range and probability distribution of the random setup error is the same along the two axes of motion.

As described in Section 5.3.1, the parameter γ is used to adjust the probabilities of each point in the support of the distribution. Values of γ that correspond to several values for the variance σ^2 of the shifts are chosen to illustrate the benefit of spatiotemporal fractionation across a spectrum of probability distributions. (The variance is measured in units of voxels instead of mm.) In the distributions of 1-voxel shifts, the variance ranges from 0 to a maximum value of $2/3$, and for the 2-voxel distributions, the variance ranges from 0 to 2. This corresponds to nominal optimization when the variance is 0 to a uniform probability distribution of the shifts when the variance is maximized.

5.4.3 Optimizing robust uniformly fractionated reference plans

Uniformly fractionated treatment plans with 21 equispaced coplanar beams were optimized to approximate a VMAT plan. Because we considered two-dimensional patient slices, we only optimized the beamlet weights of a single row of a $1 \text{ cm} \times 1 \text{ cm}$ grid of beamlets. Uniformly fractionated treatment plans with were optimized for six values of σ^2 for 2-voxel shifts and five values of σ^2 for 1-voxel shifts. The optimization problem (5.4) in Section 5.2.2 was solved to global optimality using the solver L-BFGS-B [112]. As detailed in Section 5.3.2, the quasi-Monte Carlo approximations for the expected GTV penalty are not necessary because the exact penalty values can be computed by taking advantage of the symmetry of the fractions. Optimization runtimes ranged from a few minutes with 1-voxel shifts to 28 hours for 2-voxel shifts.

The uniformly fractionated treatment plans served as high-quality reference plans that were used as a point of comparison for the spatiotemporally fractionated plans optimized in the next

section. To ensure that the uniformly fractionated plans were sufficiently robust, the penalty weights q_i in the objective function of (5.4) were adjusted so that in each of the cases, there was a 95% probability that 95% of the GTV received the prescribed 100 Gy BED.

5.4.4 Optimizing robust spatiotemporal plans

After the uniformly fractionated reference plans were computed, we solved the optimization problem (5.5) in Section 5.2.3 to compute robust spatiotemporal treatment plans. Because of the large number of scenarios in the spatiotemporal models, we used quasi-Monte Carlo sampling to approximate the expected GTV penalty with a scenario sample $\tilde{\Omega}$, as described in Section 5.3.2.

Locally optimal spatiotemporal plans were obtained by using an in-house implementation of the augmented Lagrangian algorithm. For $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for $i = 1, \dots, m$, the general inequality-constrained optimization problem

$$\begin{aligned} \min_x \quad & f(x) \\ \text{s.t.} \quad & g_i(x) \leq 0 \quad \forall i = 1, \dots, m \end{aligned} \tag{5.13}$$

can be solved with the augmented Lagrangian method (also known as the method of multipliers) [46, 74]. This method solves a sequence of unconstrained optimization problems in which the augmented Lagrangian function

$$\mathcal{L}_\rho(x, \mu) = f(x) + \frac{1}{2\rho} \sum_{i=1}^m \left[(\mu_i + \rho g_i(x))_+^2 - \mu_i^2 \right] \tag{5.14}$$

is minimized with respect to x , with $\mu = (\mu_1, \dots, \mu_m) \in \mathbb{R}_+^m$ a vector of Lagrange multipliers for the constraints $g_i(x) \leq 0$ [76]. Starting with an arbitrary $\mu^1 \geq 0$, in iteration k of the algorithm, we obtain

$$x^k = \arg \min_x \mathcal{L}_\rho(x, \mu^k) \tag{5.15}$$

for a fixed value of μ^k , and then the vector of multipliers is updated by

$$\mu^{k+1} = \left(\mu^k + \rho g(x^k) \right)_+. \tag{5.16}$$

The penalty parameter $\rho > 0$ may remain constant throughout the algorithm, or in some variations of the algorithm, ρ is increased in each iteration to improve the convergence rate of the multipliers to their optimal values [7]. The iterations continue until, for some tolerances $\epsilon_1 > 0$ and $\epsilon_2 > 0$, the stopping criteria $\|\nabla \mathcal{L}_\rho(x^k, \mu^k)\|_\infty \leq \epsilon_1$ and $\|(g(x^k))_+\|_\infty \leq \epsilon_2$ are satisfied.

In our numerical experiments, we used the augmented Lagrangian algorithm to solve the constrained optimization problem (5.5) using penalty functions of the form (5.3) for each clinical objective. All of the penalties F_i are functions of the beamlet weights x after eliminating b and d from the formulation. The function F_1 is the mean BED penalty for the uninvolved liver, which plays the role of f in the general optimization problem (5.13). The function F_2 is the penalty function to penalize BED in the GTV less than 100 Gy, and F_3 is the penalty function for BED in the unclassified tissue. We computed the expected value of penalty F_2 and the nominal penalty value of F_3 for the uniform reference plan, denoted by $\tilde{F}_2(b^*)$ and $F_3(b^*)$, respectively. Then the constraints g_1 and g_2 in the general optimization problem (5.13) are $\mathbb{E}[F_2(b)] - \tilde{F}_2(b) \leq 0$ and $F_3(b) - F_3(b^*) \leq 0$, with corresponding Lagrangian multipliers $\mu_1 \in \mathbb{R}_+$ and $\mu_2 \in \mathbb{R}_+$. The vector of multipliers $\bar{\mu} \in \mathbb{R}_+^B$ is for the nonnegativity constraint $x \geq 0$. In iteration k , we fix the multiplier values μ_1^k , μ_2^k , and $\bar{\mu}^k$ and minimize the augmented Lagrangian function

$$\begin{aligned} \mathcal{L}_\rho(x, \mu_1^k, \mu_2^k, \bar{\mu}^k) = & \frac{1}{|\tilde{\Omega}|} \sum_{\omega \in \tilde{\Omega}} F_1(b^\omega) + \frac{1}{2\rho} \left[\left(\mu_1^k + \rho \left(\left(\frac{1}{|\tilde{\Omega}|} \sum_{\omega \in \tilde{\Omega}} F_2(b^\omega) \right) - \tilde{F}_2(b^*) \right) \right)_+^2 - (\mu_1^k)^2 \right] \\ & + \frac{1}{2\rho} \left[\left(\mu_2^k + \rho (F_3(b) - F_3(b^*)) \right)_+^2 - (\mu_2^k)^2 \right] + \frac{1}{2\rho} \left[\left(\bar{\mu}^k - \rho x \right)^2 - (\bar{\mu}^k)^2 \right]. \end{aligned} \quad (5.17)$$

Once the minimum is attained, the multipliers are updated as in (5.16). (The extra clinical objectives in Cases 4 and 5 add similar terms to the augmented Lagrangian function, which we omit here for cleaner notation.)

We implemented two modifications to the standard augmented Lagrangian method described above. First, we implemented an inexact minimization of the inner problem (5.15). Previous work has proven convergence of the augmented Lagrangian method when the stopping criterion of the

inner problem (5.15) is

$$\left\| \nabla \mathcal{L}_\rho(x^k, \mu_1^k, \mu_2^k, \bar{\mu}^k) \right\|_\infty \leq \epsilon_k \quad (5.18)$$

for a nonnegative sequence $\{\epsilon_k\} \rightarrow 0$. Rockafellar proved the convergence of this inexact minimization with convex, not-necessarily-differentiable objective and gradient functions f and g_i as long as the sequence $\{\epsilon_k\}$ satisfies $\sum_{k=1}^{\infty} \sqrt{\epsilon_k} < \infty$ [77]. Bertsekas proved the convergence of the method with the inexact minimization criterion (5.18) for twice continuously differentiable f and g_i with *any* sequence $\{\epsilon_k\}$ that converges to zero, but it was noted that the inexact minimization could lead to a substantial degradation in the overall convergence of the algorithm if $\{\epsilon_k\}$ converges to zero slowly [7].

Others have demonstrated numerically [59] and proven [72] the convergence of a variation of the augmented Lagrangian method with a single gradient step towards the minimizer of $\mathcal{L}(x, \mu^k)$ in the case where f and all g_i have continuous first and second partial derivatives. In the current application, we found that limiting the number of inner iterations to 100 and allowing the inner minimization to terminate sooner if the stopping criteria are satisfied resulted in convergence of the algorithm. This heuristic increased the number of primal iterations where the minimization (5.15) is performed, but it reduced the total number of overall function evaluations compared to trials where exact minimization was used. With exact minimization, the algorithm uses a large number of iterations (and consequently, function evaluations) to solve the problem (5.15) with multipliers that are potentially very far from their optimal values. When an upper bound on the number of inner iterations is implemented, the minimization of the augmented Lagrangian is only solved approximately with the potentially inaccurate multiplier values, and as the multipliers converge to their optimal values, the inner minimizations are solved to greater accuracy.

The second variation in our implementation of the augmented Lagrangian method was in the evaluation of the constraint violation for the multiplier update (5.16). For the expected penalty terms, we evaluated the constraint violation on a larger lattice sample than what was used in the inner iterations for the minimization (5.15). This evaluation would be too time-consuming to carry out in each inner iteration of the algorithm because of the expensive gradient evaluation required in

each iteration, but evaluating the constraint violation during the multiplier update does not require the expensive gradient calculation. The large lattice size allows for greater approximation accuracy in the multiplier update.

5.4.5 Computing equieffective dose

In Chapter 3, we computed the equivalent dose DEQ5 of the BED for an interpretable comparison between spatiotemporally fractionated treatment plans and their uniformly fractionated reference plans. In this chapter, we compute the DEQ5 of the *expected* BED, which is the physical dose distribution that if delivered in 5 uniform fractions would yield the same BED as the expected BED for each voxel. The DEQ5 of the expected BED is the quantity satisfying

$$\mathbb{E} \left[5d \left(1 + \frac{d}{\alpha/\beta} \right) \right] = \text{DEQ5} + \frac{(\text{DEQ5})^2/5}{\alpha/\beta}$$

for uniformly fractionated plans and

$$\mathbb{E} \left[\sum_{t=1}^5 \left(d_t + \frac{d_t^2}{\alpha/\beta} \right) \right] = \text{DEQ5} + \frac{(\text{DEQ5})^2/5}{\alpha/\beta}$$

for spatiotemporal plans.

5.5 Results for clinical liver cases

The benefit of spatiotemporal fractionation is measured by the reduction in expected mean liver BED in the healthy liver tissue from the uniform reference plan. By construction, the spatiotemporal plans have the same plan quality as the uniform reference plan with respect to all other clinical objectives. Tables 5.1, 5.2, and 5.3 summarize the mean liver BED reductions in spatiotemporal plans for all of the cases and probability distributions in our numerical experiments. Table 5.1 contains the reductions for probability distributions supported on shifts up to 2 voxels (approximately 5 mm) in each axis, and Tables 5.2 and 5.3 display the reductions in mean liver BED

Table 5.1: Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 2-voxel shifts in either axis of the transverse plane. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to 2, where there is an equal probability of shifting to the 25 points around the nominal position. The voxel sizes are 2.54 mm \times 2.54 mm for Cases 1-4 and 2.18 mm \times 2.18 mm for Case 5.

Case description		σ^2					
		0	$\frac{1}{4}$	$\frac{1}{2}$	1	$\frac{3}{2}$	2
1	Large central lesion	26.8	8.3	5.5	4.1	3.0	2.7
2	Small lesion	32.4	13.6	12.3	8.8	7.9	6.9
3	Two small lesions	28.6	11.6	9.4	7.5	6.0	4.6
4	Lesion abutting ch. wall	50.1	24.9	19.2	14.1	11.6	10.7
5	Lesion abutting GI tract	23.6	4.9	3.0	1.1	1.1	1.0

Table 5.2: Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 1-voxel shifts in either axis of the transverse plane. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to $2/3$, where there is an equal probability of shifting to the 9 points around the nominal position. The voxel sizes are 2.54 mm \times 2.54 mm for Cases 1-4 and 2.18 mm \times 2.18 mm for Case 5.

Case description		σ^2				
		0	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{2}{3}$
1	Large central lesion	26.8	10.1	8.3	6.8	6.1
2	Small lesion	32.4	16.6	16.4	15.8	15.1
3	Two small lesions	28.6	14.5	12.9	12.0	11.2
4	Lesion abutting ch. wall	50.1	29.9	25.7	23.5	21.0
5	Lesion abutting GI tract	23.6	7.7	5.4	4.2	3.2

Table 5.3: Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 1-voxel shifts in either axis of the transverse plane with the full-resolution dose grid. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to $2/3$, where there is an equal probability of shifting to the 9 points around the nominal position. The voxel sizes are $1.27 \text{ mm} \times 1.27 \text{ mm}$ for Cases 1-4 and $1.09 \text{ mm} \times 1.09 \text{ mm}$ for Case 5.

Case description		σ^2				
		0	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{2}{3}$
1	Large central lesion	25.6	17.5	15.0	13.2	12.4
2	Small lesion	34.4	28.5	26.0	24.6	23.6
3	Two small lesions	28.6	22.8	20.7	19.4	18.6
4	Lesion abutting ch. wall	50.0	43.0	39.8	37.6	35.9
5	Lesion abutting GI tract	25.5	17.6	14.4	12.7	11.8

of spatiotemporal plans for probability distributions supported on 1-voxel shifts for downsampled (approximately 2.5 mm) and full-resolution (approximately 1.25 mm) dose grids. For reference, the tables also report the mean liver BED reductions that could be obtained in the case that no uncertainty is present. Note that the reductions for the $\sigma^2 = 0$ distributions are larger than those reported in Chapter 3; this is because the treatment plans in Chapter 3 had a PTV surrounding the GTV, whereas the treatments in this chapter are optimized without a PTV.

As expected, the benefit of spatiotemporal fractionation decreases with larger setup uncertainty. However, when the variance in the patient position is small, spatiotemporally fractionated plans maintain a large reduction in expected mean liver BED. With the downsampled dose grid, for the smallest nonzero σ^2 , Case 4 displays the largest reduction of 24.9% for 2-voxel shifts, while Case 5 has the smallest reduction of 4.9%. The former corresponds to a DEQ5 reduction from 19.3 to 16.6 Gy, while the latter corresponds to a DEQ5 reduction from 33.0 to 32.1 Gy. The reductions are even larger for spatiotemporal plans with 1-voxel shifts on both the downsampled and the full-resolution dose grid. This suggests that the magnitude of the largest undetected random patient setup error is just as important as the variance of the setup error.

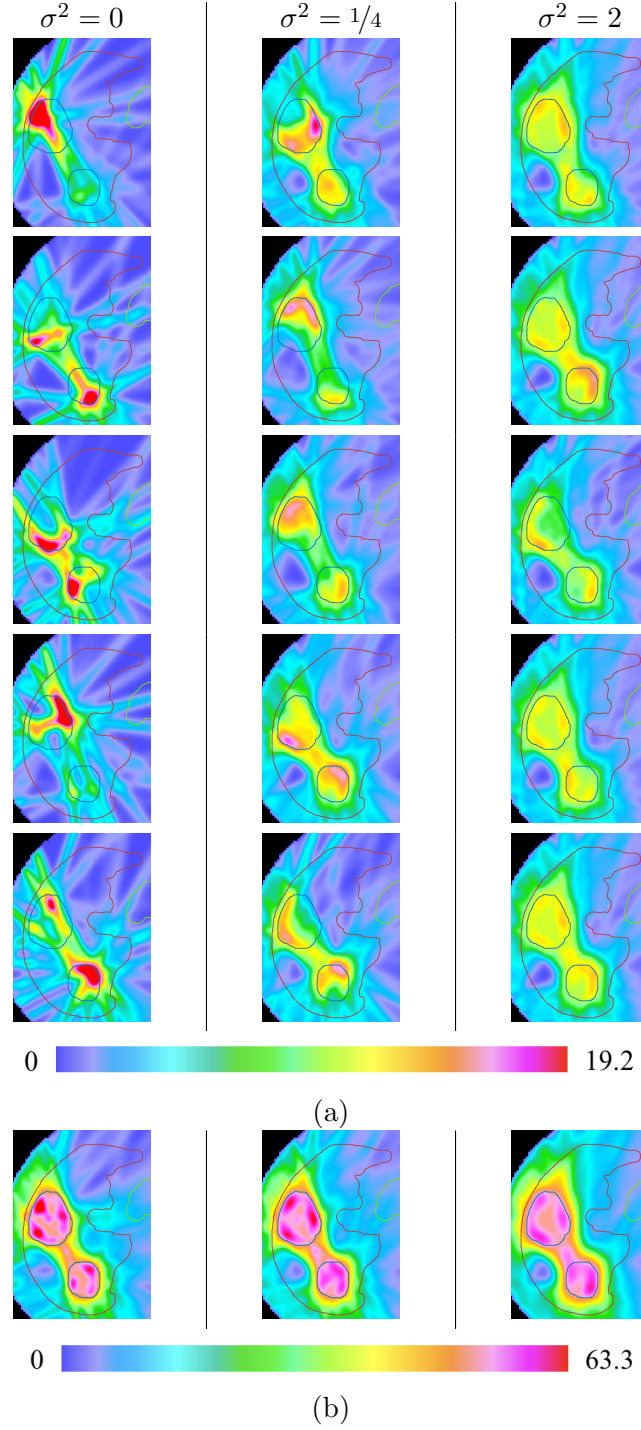


Figure 5.1: (a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 3 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.

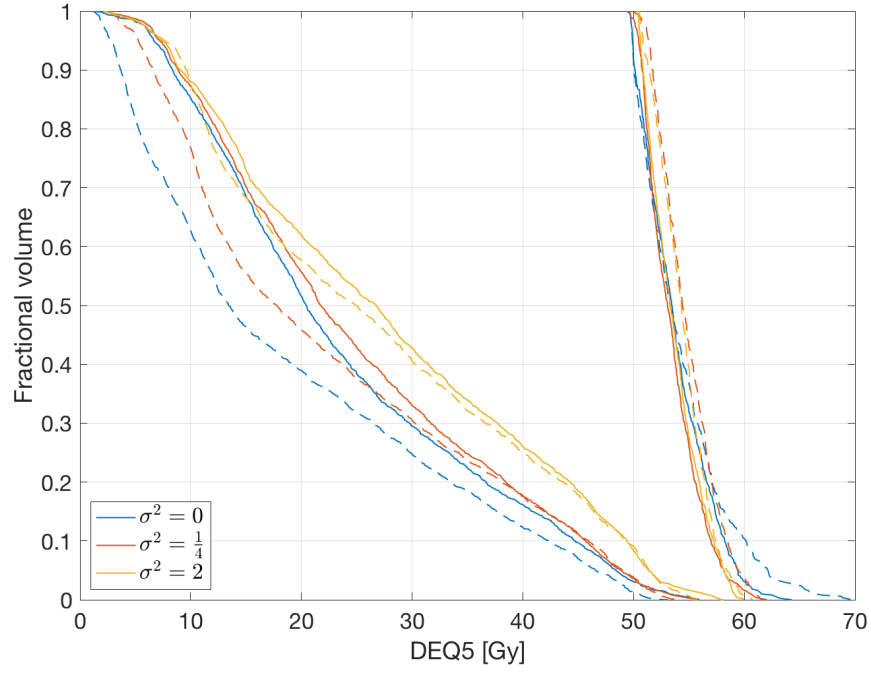
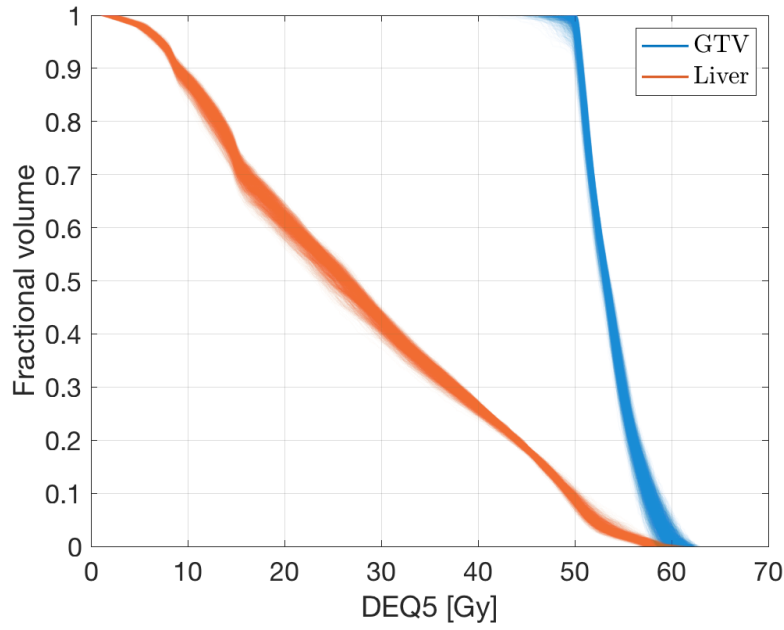
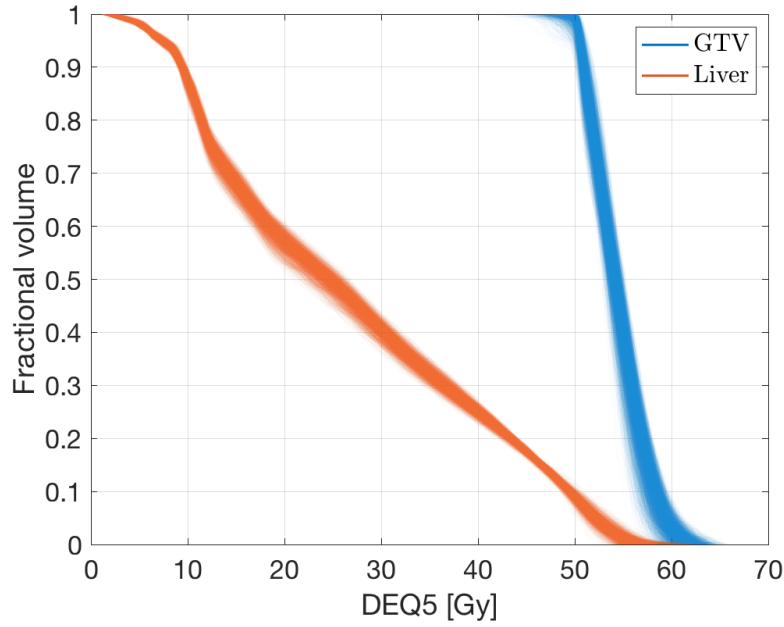


Figure 5.2: Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 3, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.



(a)



(b)

Figure 5.3: A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 3 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.

As the variance σ^2 increases, the fractions of the spatiotemporal plans become increasingly uniform. As seen in Figure 5.1a, the fractional doses are highly modulated for the smaller variances, i.e. each fraction delivers a very high dose to complementary parts of the tumor. As the variance increases, the fractional dose distributions become more uniform, i.e. each fraction delivers a significant dose to most of the target volume. This trend is expected because more uniform dose distributions help maintain robustness against greater misalignment of the fractional dose distributions. This explains the decreasing benefit of spatiotemporal fractionation over uniform reference plans as σ^2 increases.

In Figure 5.1b, the DEQ5 of the expected BED for Case 3 is plotted for the spatiotemporal plan for three σ^2 values. The robust treatment plans exhibit a “margin” of higher dose around the GTV to maintain robustness despite the fact that there is no explicit margin, such as the PTV, contoured around the GTV. As seen in the figure, the width of the “margin” around the GTV increases as the variance σ^2 of patient shifts increases. The stochastic optimization model adjusts the width of the margin around the GTV depending on the probabilities of patient shifts, and this margin is optimized to exactly maintain robustness but also to avoid unduly irradiating the healthy liver with too large a margin. These trends in the dose distributions were observed in all of the five cases; see additional figures in Appendix B.

Figure 5.2 contains the dose-volume histogram (DVH) curves for Case 3, which also shows the evolution of the optimal treatment plans as the variance of patient shifts increases. In this figure, we observe that as σ^2 increases, the healthy liver sparing deteriorates, and the curves for the uniform and spatiotemporal plans move closer to one another.

A key feature of the robust spatiotemporal plans is that the reduction in mean BED in the healthy liver tissue occurs simultaneously with equally robust coverage of the tumor. In all of the cases presented in this chapter, the spatiotemporal plan meets the same robustness criterion as its uniform reference plan: there is at least a 95% probability that 95% of the GTV receives 100 Gy BED. In most cases we found that even though the spatiotemporal plan and the uniform plan have equal expected GTV penalties, the spatiotemporal plans generally have a higher probability that

95% of the GTV receives the prescription 100 Gy BED. In addition, in most cases the spatiotemporal plan has a higher expected mean GTV BED than the uniformly fractionated reference plan.

The similar robustness between the two fractionation schemes can be seen in Figure 5.3, which displays a DVH plot for the DEQ5 values comparing the uniform reference plan and the spatiotemporally fractionated treatment plan for Case 3. Figure 5.3 displays a “cloud” of DVH curves that shows the DVH curve for every scenario in a sample with 4096 scenarios. The spatiotemporal plan exhibits more variability in the tumor, as evidenced through a greater spread of DVH curves, but we notice that the variability occurs in the region of higher DEQ5. The spatiotemporal plan maintains the same robust GTV coverage as the uniform reference plan in the critical region around DEQ5 of 50 Gy. DVH plots similar to Figure 5.2 and Figure 5.3 for the remaining cases can be found in Appendix B.

5.6 Discussion

Evaluating the benefit. In this chapter we have demonstrated that spatiotemporal fractionation schemes retain a portion of their benefit in healthy tissue sparing even in the presence of random patient setup uncertainty. As expected, the benefit of spatiotemporal fractionation decreases with increasing setup uncertainty. However, even with patient positioning errors of up to 5 mm along both axes, spatiotemporal fractionation schemes yielded a reduction in the expected mean liver BED over uniformly fractionated plans in each case, without any compromise in robustness or in the other clinical objectives. For smaller setup errors of 1-2 mm, which is applicable to intracranial lesions, a large portion of the benefit remains.

Measures of robustness. In our optimization models, the expected value of the piecewise quadratic penalty of underdose in the GTV is used as the primary measure of the robustness of a treatment plan. The constraints in the spatiotemporal optimization model are active at the solution, which means that the expected penalty values for the GTV are equal in the uniform and spatiotemporal treatment plans. Thus, with respect to the expected GTV penalty robustness mea-

sure, the uniformly fractionated treatment plans and the spatiotemporal plans are equally robust. In practice, robustness of a treatment plan is often assessed with a DVH-like criterion, such as the probability that 95% of the GTV receives the prescription BED [42, 99]. An exact optimization model with a DVH-based robustness measure would require a very large mixed-integer program with binary variables for all of the GTV voxels in every scenario, which is computationally intractable. In our models we used the expected GTV penalty to quantify the robustness, but after the optimization we also evaluated the robustness of the optimal solutions using DVH criteria. We have found that in nearly every experiment, the spatiotemporal plans are more robust than the uniform reference plans with respect to this robustness measure. Thus, the healthy tissue sparing could potentially be even higher than reported if the spatiotemporal plans and the uniform plans were equally robust with respect to the DVH-based robustness criterion.

Other sources of uncertainty. In this work, we focus on random setup uncertainty. Setup errors that are different between fractions yield a misalignment of BED contributions from different fractions, and hence represent the main concern in spatiotemporal fractionation. In addition, random errors could not be accounted for via margins because the dose distributions per fraction are determined during the optimization and not beforehand. In future work, systematic errors should also be included in robust treatment plan optimization. However, systematic errors lead to a shift (rather than a degradation) of the cumulative BED distribution. Therefore, systematic errors affect spatiotemporal treatments in the same way as conventional treatments, and could, in principle, be mitigated via PTV margins. For the application to liver tumors, intrafraction motion of the tumor due to respiration represents an additional uncertainty. However, modern treatment devices such as MR-linacs, which allow for gated treatment based on real-time MR imaging, may be used to mitigate such uncertainties.

Chapter 6

Conclusions and open problems

In this thesis, we have applied mathematical optimization techniques to improve treatment planning optimization in photon arc radiotherapy. We have achieved a better trade-off of treatment time versus treatment quality for arc therapy plans by using a greedy merging strategy and by including an upper bound on the treatment time in the fluence map optimization. We have shown that the optimization problem with an explicit upper bound on the treatment time has a special structure that allows the optimal treatment plan to be computed efficiently. Further, we studied the efficacy of a treatment modality in photon arc therapy called spatiotemporal fractionation, which allows different dose distributions to be delivered on each treatment day. We demonstrated the benefit of spatiotemporal fractionation in clinical liver cases, mathematically derived an upper bound on the maximum achievable benefit, and showed that spatiotemporal plans attain a large percentage of the maximum benefit. We also addressed the problem of patient setup uncertainty in spatiotemporally fractionated treatments by developing a stochastic optimization model to account for the uncertainties in patient positioning, and we optimized spatiotemporal plans that are robust against these uncertainties.

This thesis extends the work of radiotherapy treatment planning optimization, which has a relatively recent history. The introduction of IMRT in the early 1990s by Bortfeld and Webb [13,

102] was transformative in planning radiotherapy treatments that achieve better dose conformality around target structures, and IMRT treatments have attained widespread use in the clinic since the early 2000s [12]. Spatiotemporal fractionation is a more recent development; it was first proposed for proton radiotherapy in the early 2010s [97, 98, 110], and a more recent proof-of-concept paper extended the idea of spatiotemporal fractionation for photon radiotherapy [89]. The material in this thesis is part of the early work demonstrating the benefit of spatiotemporal plans by computing photon treatment plans for clinical cases. As we have shown, spatiotemporal fractionation has great potential to even further improve healthy tissue sparing in photon radiotherapy treatments by delivering different dose distributions in each fraction.

One direction for future research is the development of algorithms to solve the large-scale optimization problems for a three-dimensional patient. In particular, the development of efficient algorithms is necessary to solve the large-scale stochastic constrained optimization problems from Chapter 5, perhaps by employing dynamic sampling techniques or an alternative model that incorporates both the expected value and variance of the BED. Further, more efficient conic optimization algorithms are needed to solve the large-scale semidefinite programming relaxations to obtain the bounds on the benefit of three-dimensional spatiotemporal plans. The results in this thesis for two-dimensional slices of clinical cases demonstrate the potential of spatiotemporal fractionation, and the future development of these algorithms for three-dimensional cases will promote the widespread clinical use of this effective treatment modality.

REFERENCES

- [1] F. Alizadeh and D. Goldfarb. Second-order cone programming. *Mathematical Programming*, 95(1):3–51, 2003.
- [2] L. K. Ballas, B. B. Elkin, D. Schrag, B. D. Minsky, and P. B. Bach. Radiation therapy facilities in the United States. *International Journal of Radiation Oncology*Biography*Physics*, 66(4):1204–1211, 2006.
- [3] M. Balvert and D. Craft. Fast approximate delivery of fluence maps for IMRT and VMAT. *Physics in Medicine & Biology*, 62(4):1225–1247, 2017.
- [4] W. A. Beckham, P. J. Keall, and J. V. Siebers. A fluence-convolution method to calculate radiation therapy dose distributions that incorporate random set-up error. *Physics in Medicine & Biology*, 47(19):3465–3473, 2002.
- [5] L. R. Bell, E. M. Pogson, P. E. Metcalfe, and L. C. Holloway. Defining and assessing an anisotropic delineation margin for modern radiotherapy. *Medical Physics*, 43(12):6644–6653, 2016.
- [6] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization*. SIAM, Philadelphia, PA, 2001.
- [7] D. P. Bertsekas. Combined primal-dual and penalty methods for constrained minimization. *SIAM Journal on Control*, 13(3):521–544, 1975.
- [8] D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*, chapter 2, pages 46–53. Athena Scientific, 1997.
- [9] A. Bertuzzi, C. Bruni, F. Papa, and C. Sinisgalli. Optimal solution for a cancer radiotherapy problem. *Journal of Mathematical Biology*, 66(1-2):311–349, 2013.

- [10] R. Bohoslavsky, M. G. Witte, T. M. Janssen, and M. van Herk. Probabilistic objective functions for margin-less IMRT planning. *Physics in Medicine & Biology*, 58(11):3563–3580, 2013.
- [11] R. Bokrantz and K. Miettinen. Projections onto the Pareto surface in multicriteria radiation therapy optimization. *Medical Physics*, 42(10):5862–5870, 2015.
- [12] T. Bortfeld. IMRT: a review and preview. *Physics in Medicine & Biology*, 51(13):R363, 2006.
- [13] T. Bortfeld, J. Bürkelbach, R. Boesecke, and W. Schlegel. Methods of image reconstruction from projections applied to conformation radiotherapy. *Physics in Medicine & Biology*, 35(10):1423–1434, 1990.
- [14] T. Bortfeld, S. B. Jiang, and E. Rietzel. Effects of motion on the total dose distribution. *Seminars in Radiation Oncology*, 14(1):41–51, 2004.
- [15] T. Bortfeld, K. Jokivarsi, M. Goitein, J. Kung, and S. B. Jiang. Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation. *Physics in Medicine & Biology*, 47(13):2203–2220, 2002.
- [16] T. Bortfeld, J. Ramakrishnan, J. Tsitsiklis, and J. Unkelbach. Optimization of radiation therapy fractionation schedules in the presence of tumor repopulation. *INFORMS Journal on Computing*, 27(4):788–803, 2015.
- [17] S. P. Boyd and L. Vandenberghe. *Convex Optimization*, chapter A5, pages 645–652. Cambridge University Press, 2004.
- [18] D. J. Brenner. The linear-quadratic model is an appropriate methodology for determining isoeffective doses at large doses per fraction. *Seminars in Radiation Oncology*, 18(4):234–239, 2008.
- [19] R. H. Byrd, M. E. Hribar, and J. Nocedal. An interior point algorithm for large-scale nonlinear programming. *SIAM Journal on Optimization*, 9(4):877–900, 1999.

- [20] T. C. Y. Chan, T. Bortfeld, and J. N. Tsitsiklis. A robust approach to IMRT optimization. *Physics in Medicine & Biology*, 51(10):2567–2583, 2006.
- [21] M. Chu, Y. Zinchenko, S. G. Henderson, and M. B. Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Physics in Medicine & Biology*, 50(23):5463–5477, 2005.
- [22] K. L. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms*, 6(4):63:1–63:30, 2010.
- [23] P. P. Connell and S. Hellman. Advances in radiotherapy and implications for the next century: A historical perspective. *Cancer Research*, 69(2):383–392, 2009.
- [24] D. Craft. Accelerating sliding window VMAT by open-out/close-in field delivery techniques. *Medical Physics*, 42(6):3705, 2015.
- [25] D. Craft, M. Bangert, T. Long, D. Papp, and J. Unkelbach. Shared data for intensity modulated radiation therapy (IMRT) optimization research: the CORT dataset. *GigaScience*, 3(1):1–12, 2014.
- [26] D. Craft, T. Halabi, H. A. Shih, and T. Bortfeld. An approach for practical multiobjective IMRT treatment planning. *International Journal of Radiation Oncology*Biology*Physics*, 69(5):1600–1607, 2007.
- [27] D. Craft, D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld. Multicriteria VMAT optimization. *Medical Physics*, 39(2):686–696, 2012.
- [28] D. Craft, D. Papp, and J. Unkelbach. Plan averaging for multicriteria navigation of sliding window IMRT and VMAT. *Medical Physics*, 41(2):021709, 2014.
- [29] D. Craft and C. Richter. Deliverable navigation for multicriteria step and shoot IMRT treatment planning. *Physics in Medicine & Biology*, 58(1):87–103, 2013.

- [30] R. Cranley and T. Patterson. Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, 13(6):904–914, 1976.
- [31] J. O. Deasy, A. I. Blanco, and V. H. Clark. CERR: A computational environment for radiotherapy research. *Medical Physics*, 30(5):979–985, 2003.
- [32] J. O. Deasy and J. F. Fowler. Radiobiology of IMRT. In A. J. Mundt and J. C. Roeske, editors, *Intensity modulated radiation therapy: A clinical perspective*, chapter 3, pages 53–74. BC Decker, 2005.
- [33] M. Ehrgott. *Multicriteria Optimization*. Springer, 2nd edition, 2005.
- [34] D. Fontanarosa, H. P. van der Laan, M. Witte, G. Shakirin, E. Roelofs, J. A. Langendijk, P. Lambin, and M. van Herk. An in silico comparison between margin-based and probabilistic target-planning approaches in head and neck cancer patients. *Radiotherapy and Oncology*, 109(3):430–436, 2013.
- [35] A. Forsgren, P. E. Gill, and M. H. Wright. Interior methods for nonlinear optimization. *SIAM Review*, 44(4):525–597, 2002.
- [36] J. F. Fowler. 21 years of biologically effective dose. *British Journal of Radiology*, 83(991):554–568, 2010.
- [37] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- [38] A. Fredriksson and R. Bokrantz. Deliverable navigation for multicriteria IMRT treatment planning by combining shared and individual apertures. *Physics in Medicine & Biology*, 58(21):7683–7697, 2013.
- [39] A. Fredriksson and R. Bokrantz. The scenario-based generalization of radiation therapy margins. *Physics in Medicine & Biology*, 61(5):2067–2082, 2016.

- [40] M. R. Gaddy and D. Papp. Improving the VMERGE treatment planning algorithm for rotational radiotherapy. *Medical Physics*, 43(7):4093–4097, 2016.
- [41] M. R. Gaddy, S. Yıldız, J. Unkelbach, and D. Papp. Optimization of spatiotemporally fractionated radiotherapy treatments with bounds on the achievable benefit. *Physics in Medicine & Biology*, 63(1):015036, 2018.
- [42] J. J. Gordon, N. Sayah, E. Weiss, and J. V. Siebers. Coverage optimized planning: Probabilistic treatment planning based on dose coverage histogram criteria. *Medical Physics*, 37(2):550–563, 2010.
- [43] D. Gries. A note on a standard strategy for developing loop invariants and loops. *Science of Computer Programming*, 2(3):207–214, 1982.
- [44] J. Guélat and P. Marcotte. Some comments on Wolfe’s ‘away step’. *Mathematical Programming*, 35(1):110–119, 1986.
- [45] E. Hall and A. Giaccia. *Radiobiology for the radiologist*, chapter 19, pages 303–326. Lippincott Williams & Wilkins, 7th edition, 2012.
- [46] M. R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4(5):303–320, 1969.
- [47] D. A. Jaffray and M. K. Gospodarowicz. Radiation therapy for cancer. In H. Gelband, P. Jha, R. Sankaranarayanan, and S. Horton, editors, *Cancer: Disease Control Priorities*, volume 3, chapter 14, pages 266–290. The International Bank for Reconstruction and Development /The World Bank, Third edition, 2015.
- [48] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. *Proceedings of the 30th International Conference on Machine Learning*, 28(1):427–435, 2013.
- [49] N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4):373–395, 1984.

- [50] H. Keller, A. Hope, G. Meier, and M. Davison. A novel dose–volume metric for optimizing therapeutic ratio through fractionation: Retrospective analysis of lung cancer treatments. *Medical Physics*, 40(8):084101, 2013.
- [51] M. Kim. *A mathematical framework for spatiotemporal optimality in radiation therapy*. PhD thesis, University of Washington, 2010.
- [52] M. Kim, A. Ghate, and M. H. Phillips. A stochastic control formalism for dynamic biologically conformal radiation therapy. *European Journal of Operational Research*, 219(3):541–556, 2012.
- [53] M. Kim and M. H. Phillips. A feasibility study of dynamic adaptive radiotherapy for nonsmall cell lung cancer. *Medical Physics*, 43(5):2153–2161, 2016.
- [54] P. L’Ecuyer and C. Lemieux. Variance reduction via lattice rules. *Management Science*, 46(9):1214–1235, 2000.
- [55] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- [56] W. Lu, M. Chen, Q. Chen, K. Ruchala, and G. Olivera. Adaptive fractionation therapy: I. Basic concept and strategy. *Physics in Medicine & Biology*, 53(19):5495–5511, 2008.
- [57] R. Marler and J. Arora. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6):369–395, 2004.
- [58] MathWorks. MATLAB optimization toolbox documentation, 2016. <https://www.mathworks.com/help/optim/ug/fmincon.html/>.
- [59] A. Miele, P. Moseley, A. Levy, and G. Coggins. On the method of multipliers for mathematical programming problems. *Journal of Optimization Theory and Applications*, 10(1):1–33, 1972.
- [60] T. Mitin and A. L. Zietman. Promise and pitfalls of heavy-particle therapy. *Journal of Clinical Oncology*, 32(26):2855–2863, 2014.

- [61] M. Mizuta, S. Takao, H. Date, N. Kishimoto, K. L. Sutherland, R. Onimaru, and H. Shirato. A mathematical study to select fractionation regimen based on physical dose distribution and the linear-quadratic model. *International Journal of Radiation Oncology*Biology*Physics*, 84(3):829–833, 2012.
- [62] M. Monz, K. Küfer, T. Bortfeld, and C. Thieke. Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*, 53(4):985–998, 2008.
- [63] MOSEK ApS. MOSEK Optimization Suite release 8.0.0.94, 2017. [http://http://docs.mosek.com/8.0/intro.pdf](http://docs.mosek.com/8.0/intro.pdf).
- [64] H. Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods*, chapter 5, pages 101–146. Society for Industrial and Applied Mathematics, 1992.
- [65] A. Niemierko. Reporting and analyzing dose distributions: A concept of equivalent uniform dose. *Medical Physics*, 24(1):103–110, 1997.
- [66] A. Niemierko and M. Goitein. Modeling of normal tissue response to radiation: The critical volume model. *International Journal of Radiation Oncology*Biology*Physics*, 25(1):135–145, 1993.
- [67] C. C. Pan, B. D. Kavanagh, L. A. Dawson, X. A. Li, S. K. Das, M. Miften, and R. K. Ten Haken. Radiation-associated liver injury. *International Journal of Radiation Oncology*Biology*Physics*, 76(3):S94–S100, 2010.
- [68] D. Papp and J. Unkelbach. Direct leaf trajectory optimization for volumetric modulated arc therapy planning with sliding window delivery. *Medical Physics*, 41(1):011701, 2014.
- [69] Particle Therapy Co-Operative Group. Particle therapy facilities in operation. <https://www.ptcog.ch/index.php/facilities-in-operation>. Accessed: 2017-11-04.

- [70] S. Petit, J. Seco, and H. Kooy. Increasing maximum tumor dose to manage range uncertainties in IMPT treatment planning. *Physics in Medicine & Biology*, 58(20):7329–7341, 2013.
- [71] D. Pflugfelder, J. J. Wilkens, and U. Oelfke. Worst case optimization: A method to account for uncertainties in the optimization of intensity modulated proton therapy. *Physics in Medicine & Biology*, 53(6):1689–1700, 2008.
- [72] B. Polyak. Methods for solving constrained extremum problems in the presence of random noise. *USSR Computational Mathematics and Mathematical Physics*, 19(1):72–81, 1979.
- [73] R. A. Popple, P. A. Balter, and C. G. Orton. Because of the advantages of rotational techniques, conventional IMRT will soon become obsolete. *Medical Physics*, 41(10):100601, 2014.
- [74] M. Powell. Algorithms for nonlinear constraints that use Lagrangian functions. *Mathematical Programming*, 14(1):224–248, 1978.
- [75] R. Reemsten and M. Alber. Continuous optimization of beamlet intensities for intensity modulated photon and proton radiotherapy. In P. Pardalos and H. E. Romeijn, editors, *Handbook of Optimization in Medicine*, pages 83–122. Springer, 2009.
- [76] R. T. Rockafellar. A dual approach to solving nonlinear programming problems by unconstrained optimization. *Mathematical Programming*, 5(1):354–373, 1973.
- [77] R. T. Rockafellar. The multiplier method of Hestenes and Powell applied to convex programming. *Journal of Optimization Theory and Applications*, 12(6):555–562, 1973.
- [78] H. E. Romeijn, J. F. Dempsey, and J. G. Li. A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine & Biology*, 49(10):1991–2013, 2004.
- [79] E. Salari, J. Unkelbach, and T. Bortfeld. A mathematical programming approach to the fractionation problem in chemoradiotherapy. *IEEE Transactions on Healthcare Systems Engineering*, 5(2):55–73, 2015.

- [80] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. Society for Industrial and Applied Mathematics, 2009.
- [81] M. Sharma, E. Weiss, and J. V. Siebers. Dose deformation-invariance in adaptive prostate radiation therapy: Implication for treatment simulations. *Radiotherapy and Oncology*, 105(2):207–213, 2012.
- [82] N. Z. Shor. Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25(6):1–11, 1987.
- [83] I. Sloan, F. Kuo, and S. Joe. Constructing randomly shifted lattice rules in weighted Sobolev spaces. *SIAM Journal on Numerical Analysis*, 40(5):1650–1655, 2002.
- [84] J. Sonke and J. Belderbos. Adaptive radiotherapy for lung cancer. *Seminars in Radiation Oncology*, 20(2):94–106, 2010.
- [85] J. F. Sturm. Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11–12(1–4):625–653, 1999.
- [86] H. Suit, S. Goldberg, A. Niemierko, A. Trofimov, J. Adams, H. Paganetti, G. T. Y. Chen, T. Bortfeld, S. Rosenthal, J. Loeffler, and T. Delaney. Proton beams to replace photon beams in radical dose treatments. *Acta Oncologica*, 42(8):800–808, 2003.
- [87] R. Svensson, P. Källman, and A. Brahme. An analytical solution for the dynamic control of multileaf collimators. *Physics in Medicine & Biology*, 39(1):37–61, 1994.
- [88] C. Thieke, T. Bortfeld, A. Niemierko, and S. Nill. From physical dose constraints to equivalent uniform dose constraints in inverse radiotherapy planning. *Medical Physics*, 30(9):2332–2339, 2003.
- [89] J. Unkelbach. Non-uniform spatiotemporal fractionation schemes in photon radiotherapy. In D. A. Jaffray, editor, *World Congress on Medical Physics and Biomedical Engineering, June 7-12, 2015, Toronto, Canada*, pages 401–404. Springer International Publishing, 2015.

- [90] J. Unkelbach, M. Alber, M. Bangert, R. Bokrantz, T. C. Chan, J. O. Deasy, A. Fredriksson, B. L. Gorissen, M. van Herk, W. Liu, H. Mahmoudzadeh, O. Nohadani, J. V. Siebers, M. Witte, and H. Xu. Robust radiotherapy planning. *Physics in Medicine & Biology*, 63(22):22TR02, 2018.
- [91] J. Unkelbach, T. Bortfeld, D. Craft, M. Alber, M. Bangert, R. Bokrantz, D. Chen, R. Li, L. Xing, C. Men, S. Nill, D. Papp, E. Romeijn, and E. Salari. Optimization approaches to volumetric modulated arc therapy planning. *Medical Physics*, 42(3):1367–1377, 2015.
- [92] J. Unkelbach, M. R. Bussière, P. H. Chapman, J. S. Loeffler, and H. A. Shih. Spatiotemporal fractionation schemes for irradiating large cerebral arteriovenous malformations. *International Journal of Radiation Oncology*Biology*Physics*, 95(3):1067–1074, 2016.
- [93] J. Unkelbach, T. C. Y. Chan, and T. Bortfeld. Accounting for range uncertainties in the optimization of intensity modulated proton therapy. *Physics in Medicine & Biology*, 52(10):2755–2773, 2007.
- [94] J. Unkelbach, D. Craft, E. Salari, J. Ramakrishnan, and T. Bortfeld. The dependence of optimal fractionation schemes on the spatial dose distribution. *Physics in Medicine & Biology*, 58(1):159–167, 2013.
- [95] J. Unkelbach and U. Oelfke. Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions. *Physics in Medicine & Biology*, 49(17):4005–4029, 2004.
- [96] J. Unkelbach and U. Oelfke. Incorporating organ movements in IMRT treatment planning for prostate cancer: Minimizing uncertainties in the inverse planning process. *Medical Physics*, 32(8):2471–2483, 2005.
- [97] J. Unkelbach and D. Papp. The emergence of nonuniform spatiotemporal fractionation schemes within the standard BED model. *Medical Physics*, 42(5):2234–2241, 2015.

- [98] J. Unkelbach, C. Zeng, and M. Engelsman. Simultaneous optimization of dose distributions and fractionation schemes in particle radiotherapy. *Medical Physics*, 40(9):091702, 2013.
- [99] M. van Herk. Errors and margins in radiotherapy. *Seminars in Radiation Oncology*, 14(1):52–64, 2004.
- [100] S. van Kranen, O. Hamming-Vrieze, A. Wolf, E. Damen, M. van Herk, and J. Sonke. Head and neck margin reduction with adaptive radiation therapy: Robustness of treatment plans against anatomy changes. *International Journal of Radiation Oncology*Biology*Physics*, 96(3):653–660, 2016.
- [101] L. Vandenberghe and S. P. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [102] S. Webb. Optimisation of conformal radiotherapy dose distribution by simulated annealing. *Physics in Medicine & Biology*, 34(10):1349–1370, 1989.
- [103] L. M. Wein, J. E. Cohen, and J. T. Wu. Dynamic optimization of a linear-quadratic model with incomplete repair and volume-dependent sensitivity and repopulation. *International Journal of Radiation Oncology*Biology*Physics*, 47(4):1073–83, 2000.
- [104] H. P. Wieser, E. Cisternas, N. Wahl, S. Ulrich, A. Stadler, H. Mescher, L. R. Müller, T. Klinge, H. Gabrys, L. Burigo, A. Mairani, S. Ecker, B. Ackermann, M. Ellerbrock, K. Parodi, O. Jäkel, and M. Bangert. Development of the open-source dose calculation and optimization toolkit matRad. *Medical Physics*, 44(6):2556–2568, 2017.
- [105] M. G. Witte, J. van der Geer, C. Schneider, J. V. Lebesque, M. Alber, and M. van Herk. IMRT optimization including random and systematic geometric errors based on the expectation of TCP and NTCP. *Medical Physics*, 34(9):3544–3555, 2007.
- [106] Q. Wu, Y. Chi, P. Y. Chen, D. J. Krauss, D. Yan, and A. Martinez. Adaptive replanning strategies accounting for shrinkage in head and neck IMRT. *International Journal of Radiation Oncology*Biology*Physics*, 75(3):924–932, 2009.

- [107] H. Yamazaki, H. Shiomi, T. Tsubokura, N. Kodani, T. Nishimura, N. Aibe, H. Uono, M. Nishikata, Y. Baba, M. Ogita, K. Yamashita, and T. Kotsuma. Quantitative assessment of inter-observer variability in target volume delineation on stereotactic radiotherapy treatment for pituitary adenoma and meningioma near optic tract. *Radiation Oncology*, 6(1), 2011.
- [108] D. Yan, F. Vicini, J. Wong, and A. Martinez. Adaptive radiation therapy. *Physics in Medicine & Biology*, 42(1):123–132, 1997.
- [109] Y. Yang and L. Xing. Optimization of radiotherapy dose-time fractionation with consideration of tumor specific biology. *Medical Physics*, 32(12):3666–3677, 2005.
- [110] C. Zeng, D. Giantsoudi, C. Grassberger, S. Goldberg, A. Niemierko, H. Paganetti, J. A. Efstathiou, and A. Trofimov. Maximizing the biological effect of proton dose delivered with scanned beams via inhomogeneous daily dose distributions. *Medical Physics*, 40(5):051708, 2013.
- [111] C. Zhu, R. Byrd, J. Nocedal, and J. L. Morales. L-BFGS-B: Software for large-scale bound-constrained optimization, version 3.0, 2011. <http://users.iems.northwestern.edu/~nocedal/lbfgsb.html>.
- [112] C. Zhu, R. H. Byrd, and J. Nocedal. L-BFGS-B, FORTRAN routines for large scale bound constrained optimization. *ACM Transactions on Mathematical Software*, 23(4):550–560, 1997.

APPENDICES

Appendix A

Supplementary figures for Chapter 3

This appendix contains the figures for the additional liver cases not presented in the main text of Chapter 3. These figures show a comparison of locally optimal spatiotemporal treatment plans and their uniformly fractionated reference plans. In each of the cases, the spatiotemporal plan achieves an improvement in healthy tissue sparing while maintaining the same tumor coverage as the uniformly fractionated treatment.

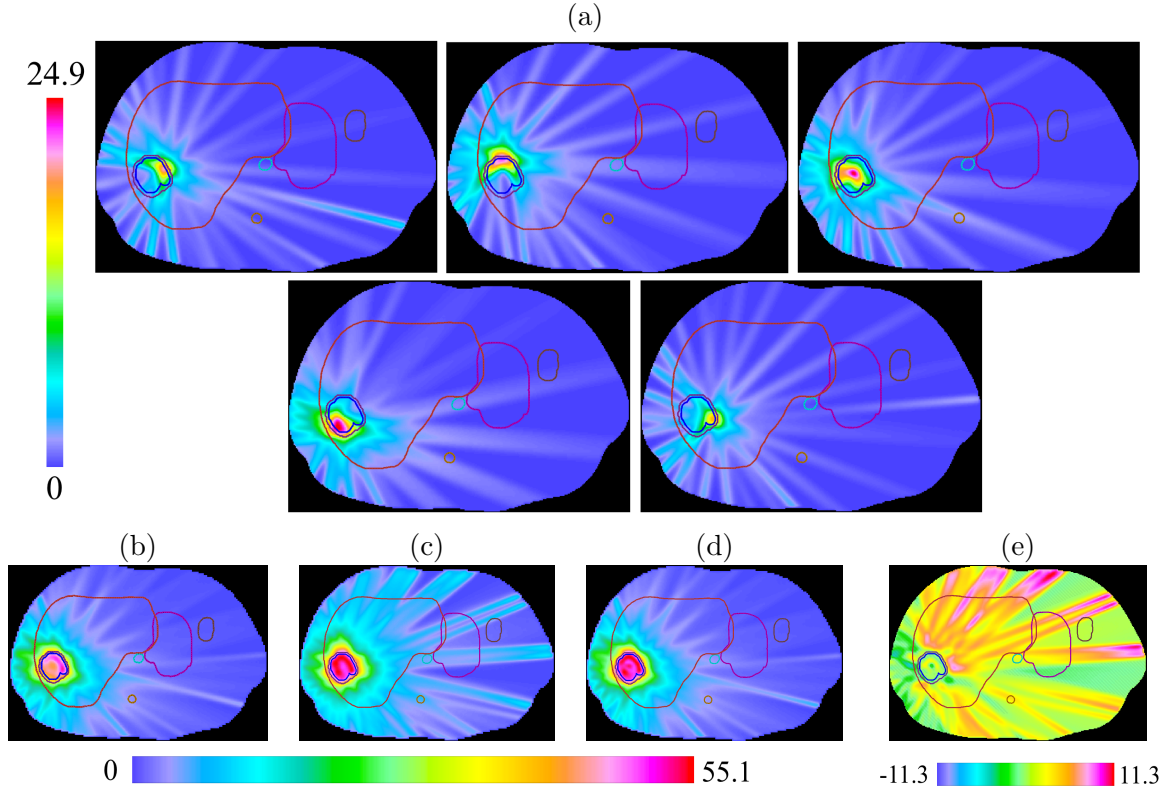


Figure A.1: Dose distributions for Case 2, a small lesion within the liver. (a) Physical dose distributions in a 5-fraction nonuniformly fractionated treatment. Although the tumor is small, the nonuniform plan still hypofractionates different subregions of the target. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan. As in Case 1, the nonuniform spatiotemporal plan provides similar target coverage as the uniform reference plan, with a more conformal dose distribution. (e) The plot of (c) minus (d) shows that, as expected, the majority of the BED reduction is in the healthy liver and in the entrance regions of the beams exposing the liver the most. All numerical quantities shown are in Gy.

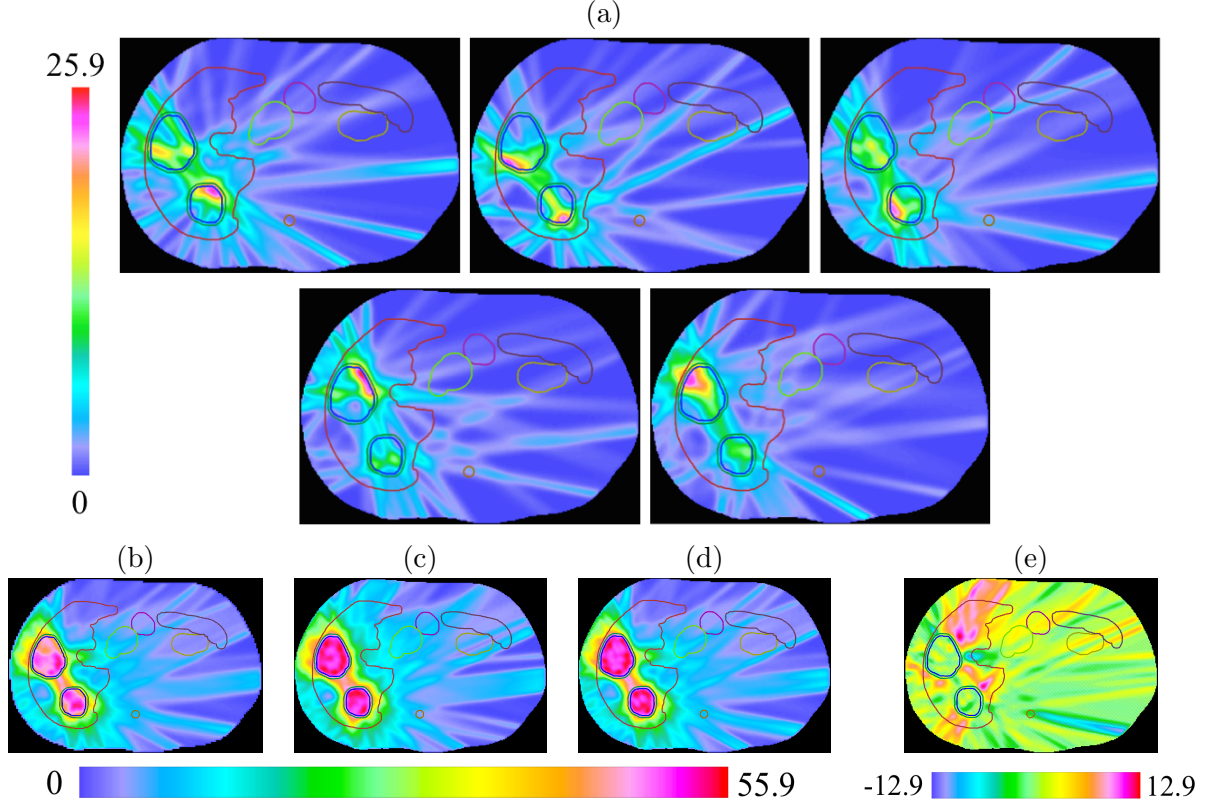


Figure A.2: Nonuniformly fractionated dose distributions for Case 3, which contains two lesions within the liver. (a) Physical dose distributions in each of the five fractions show that the nonuniformly fractionated treatment hypofractionates different parts of the tumor. (b) Total physical dose delivered throughout the nonuniformly fractionated treatment. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan, which is the uniform plan that is isoeffective in delivering the same BED as the nonuniformly fractionated plan. (e) The difference between the physical dose in the uniform plan and the DEQ5 for the nonuniform plan, or (c) minus (d). This shows that the benefit of nonuniform fractionation is mostly in reduced dose in the healthy liver and in the entrance region of the those beams that expose the liver the most. All numerical quantities shown are in Gy.

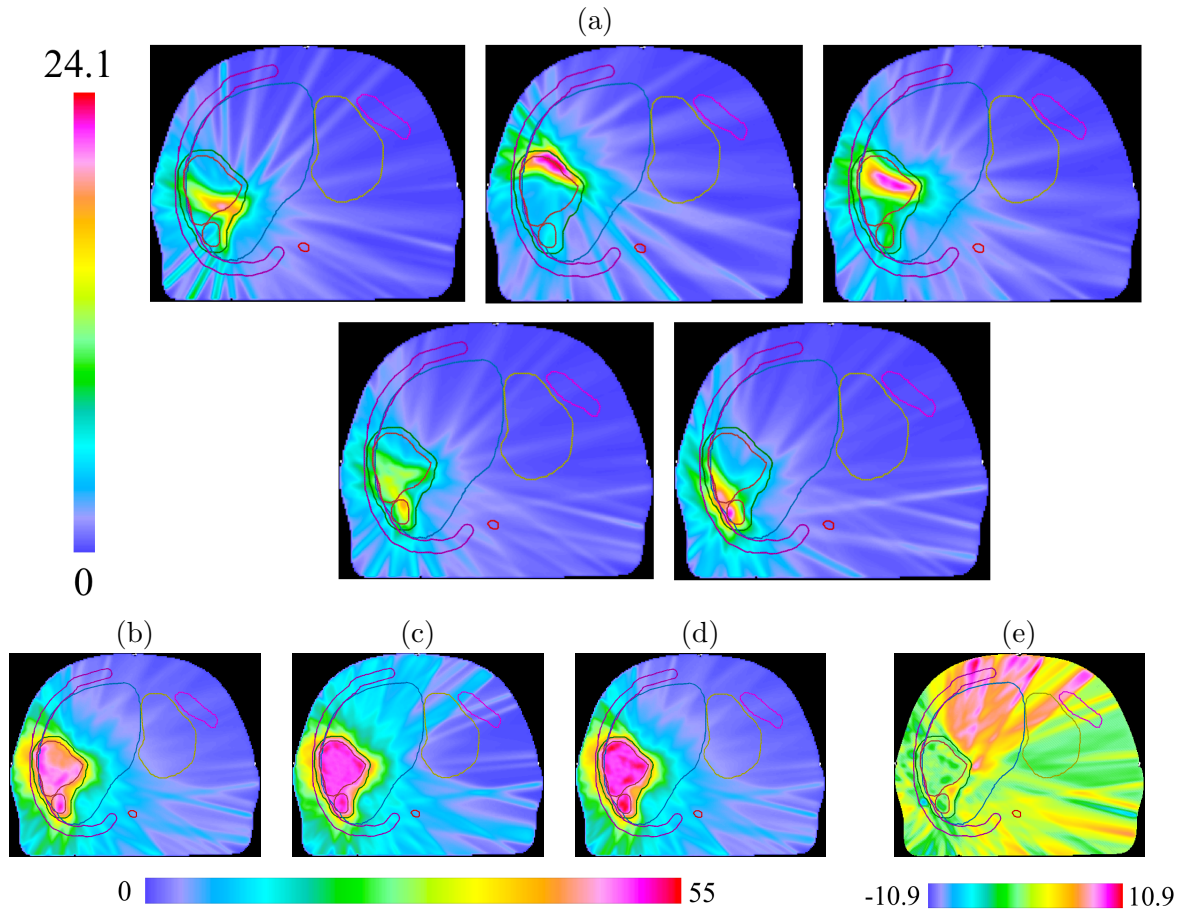


Figure A.3: Dose distributions for Case 4, a lesion abutting the chest wall. (a) Physical dose distributions for each fraction of a nonuniformly fractionated treatment. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniform reference plan. (d) DEQ5 of the nonuniformly fractionated plan. (e) Plot of (c) minus (d). All numerical quantities shown are in Gy.

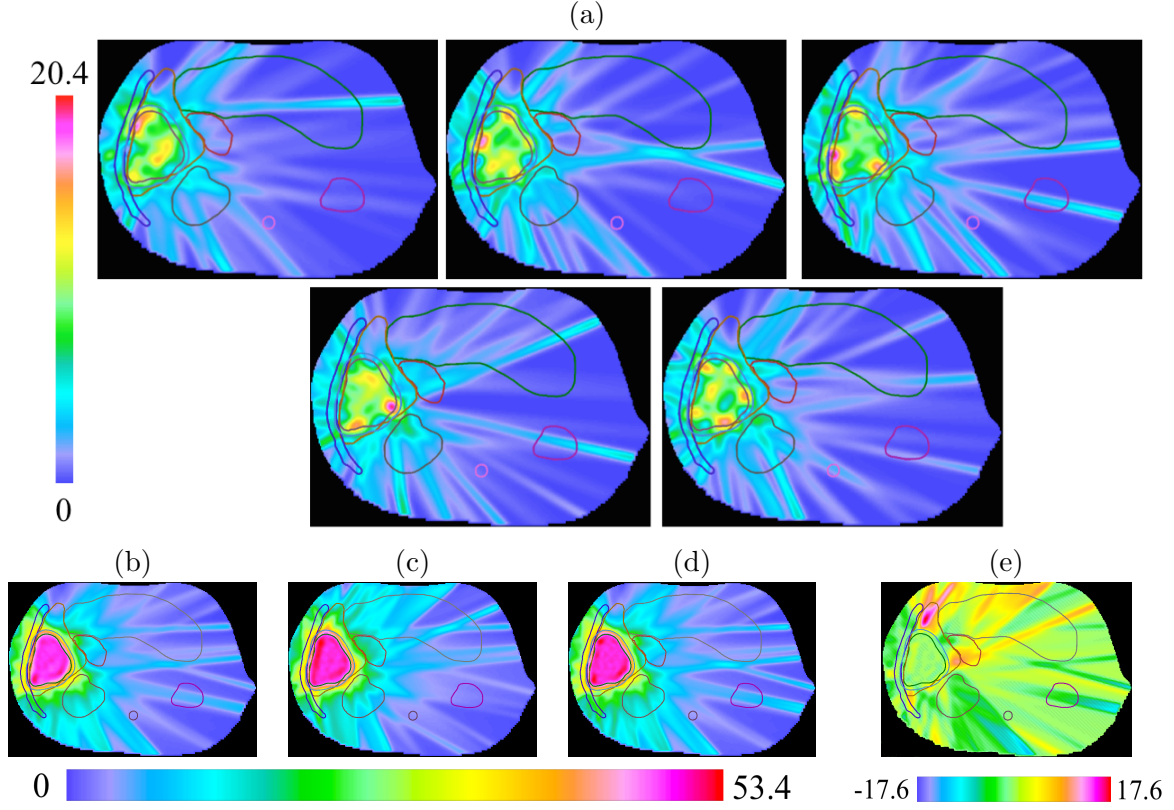


Figure A.4: Dose distributions for Case 5, a lesion abutting the GI tract. (a) Physical dose distributions for each fraction of a nonuniformly fractionated treatment. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan. (e) Plot of (c) minus (d). All numerical quantities shown are in Gy.

Appendix B

Supplementary figures for Chapter 5

This appendix contains the figures for the additional clinical cases that were not included in the main text of Chapter 5. These remaining clinical cases display the same trends as observed in Case 3 presented in the main text. As the variance σ^2 increases, the fractional dose distributions become more uniform in order to maintain robustness against larger patient setup uncertainty. From the DVH plots for each clinical case we observe that the spatiotemporal treatment plans have a reduction in healthy liver BED while maintaining the same robustness in the GTV as the uniformly fractionated plan.

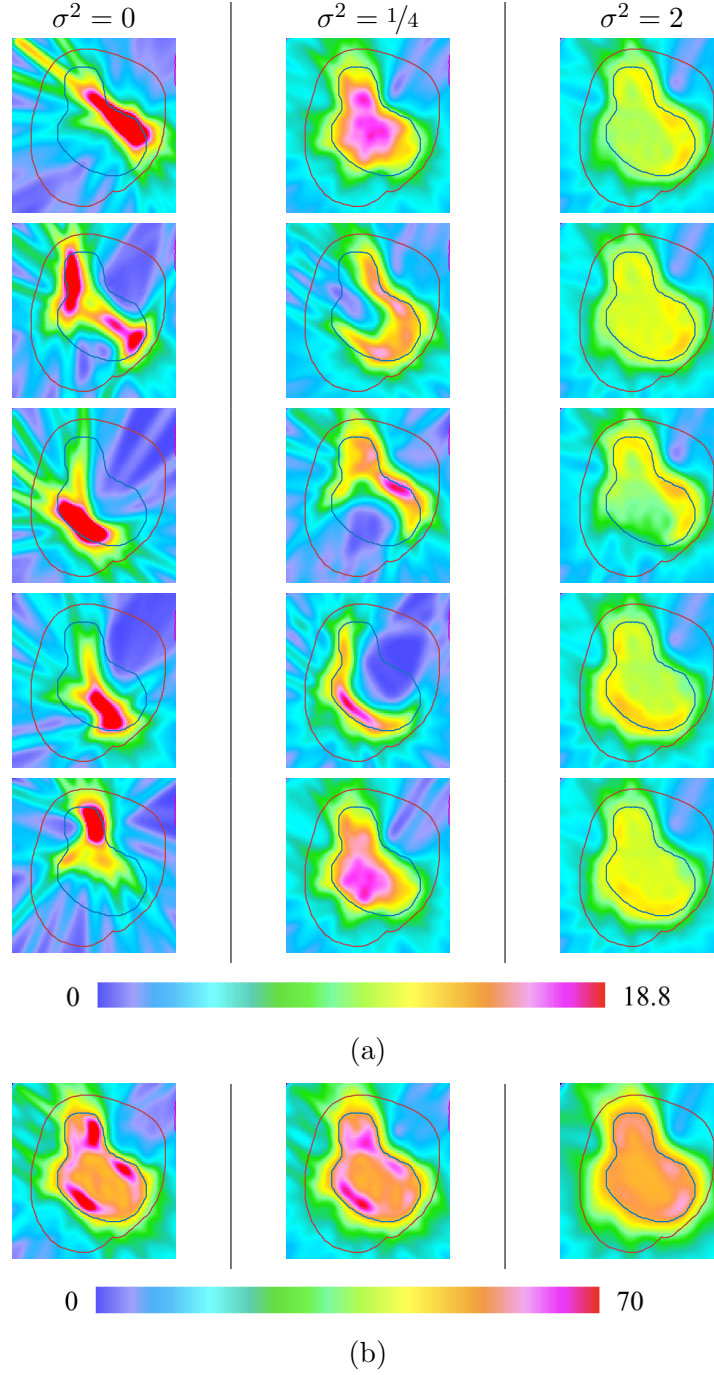


Figure B.1: (a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 1 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.

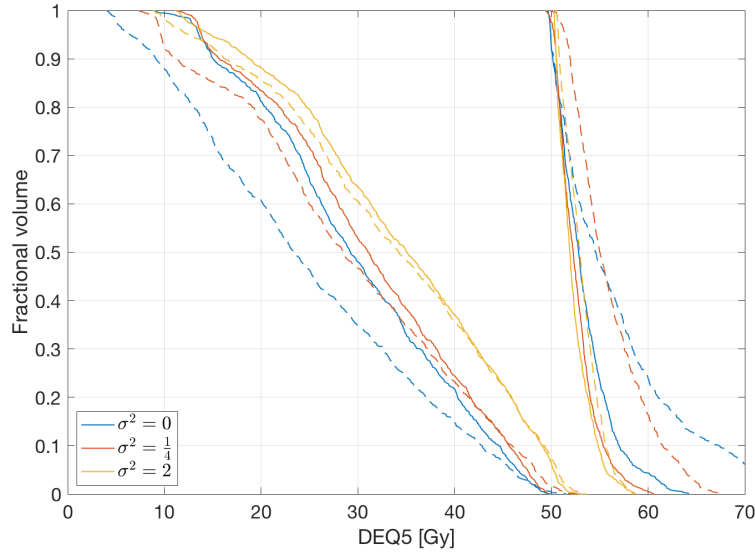


Figure B.2: Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 1, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.

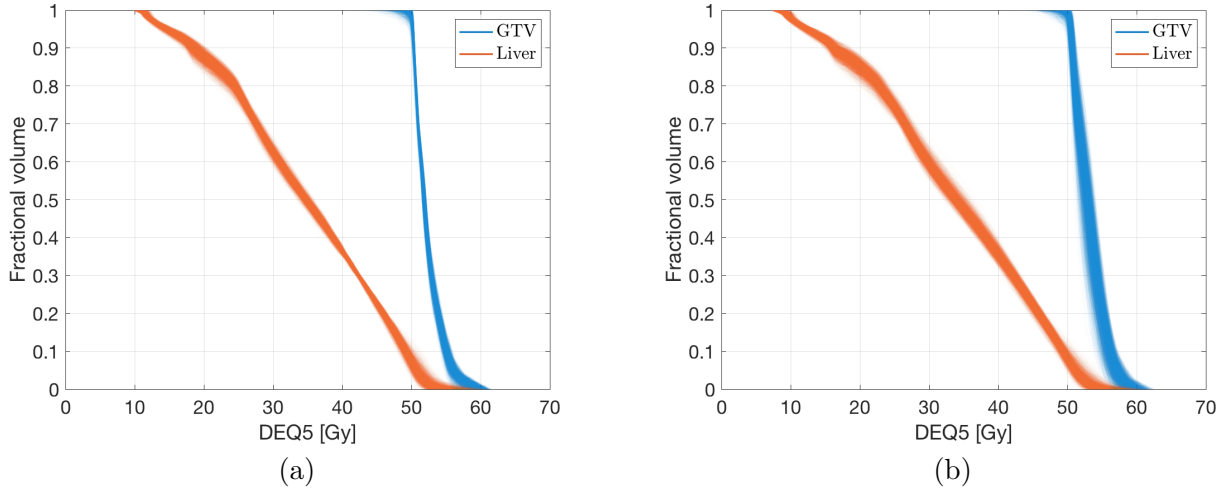


Figure B.3: A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 1 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.

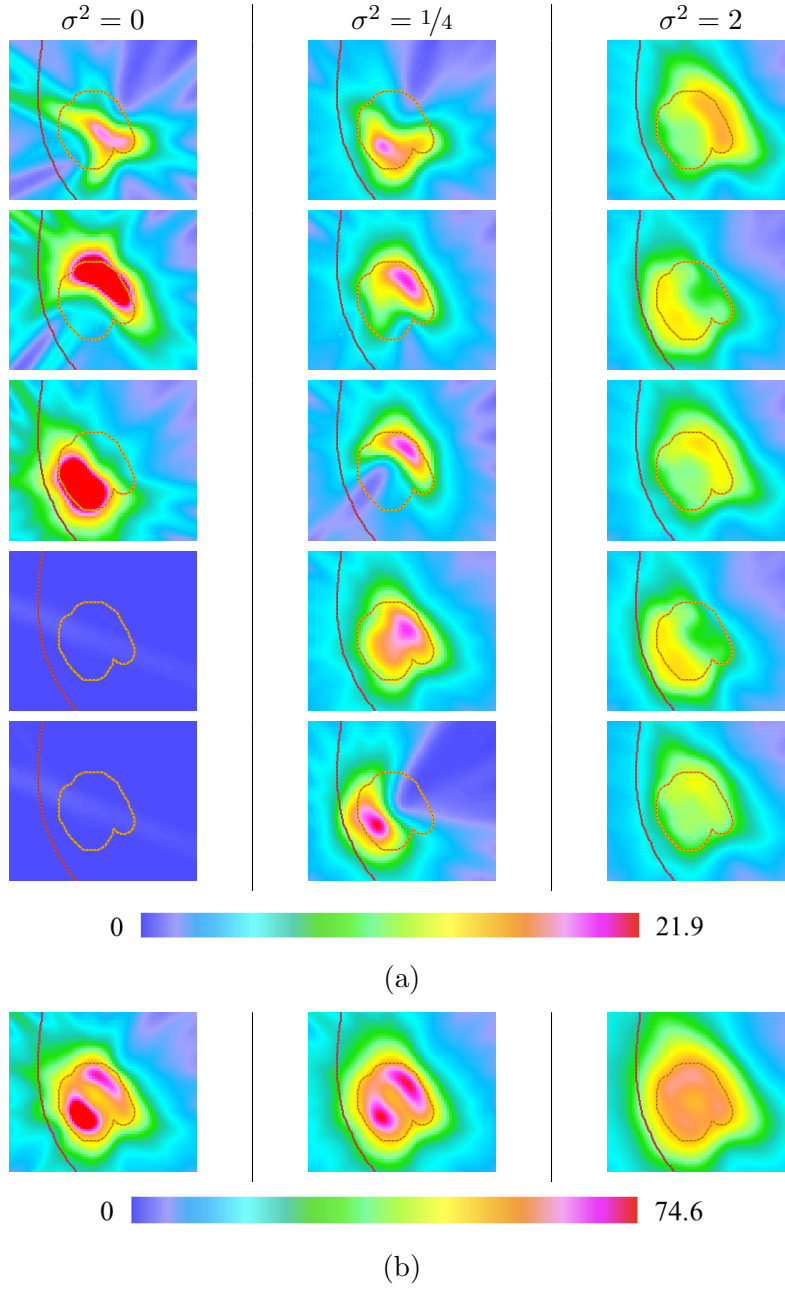


Figure B.4: (a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 2 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.

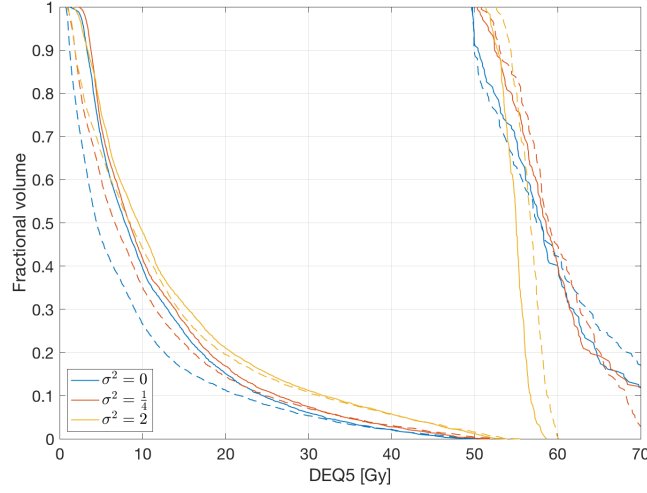


Figure B.5: Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 2, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.

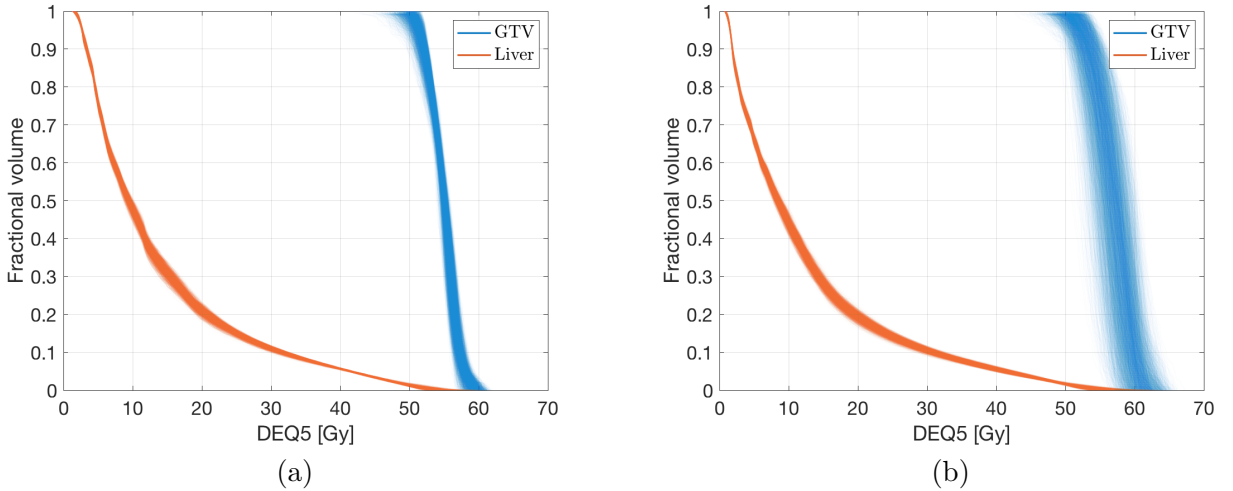


Figure B.6: A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 2 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.

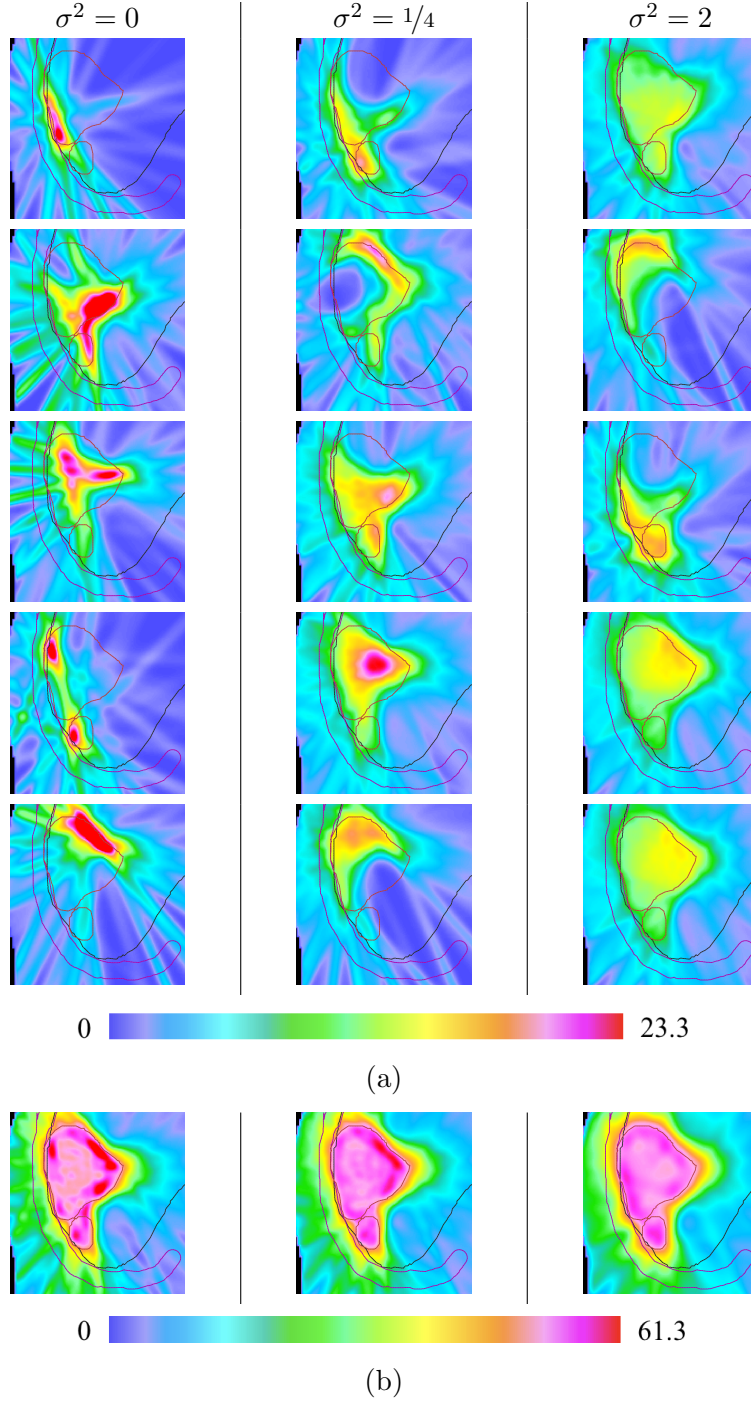


Figure B.7: (a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 4 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.

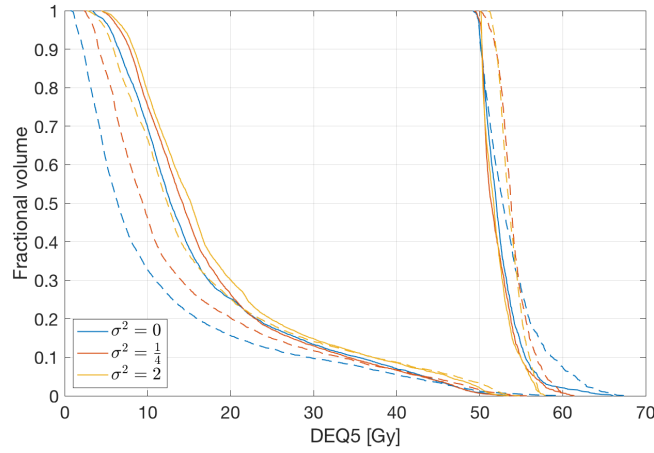


Figure B.8: Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 4, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.

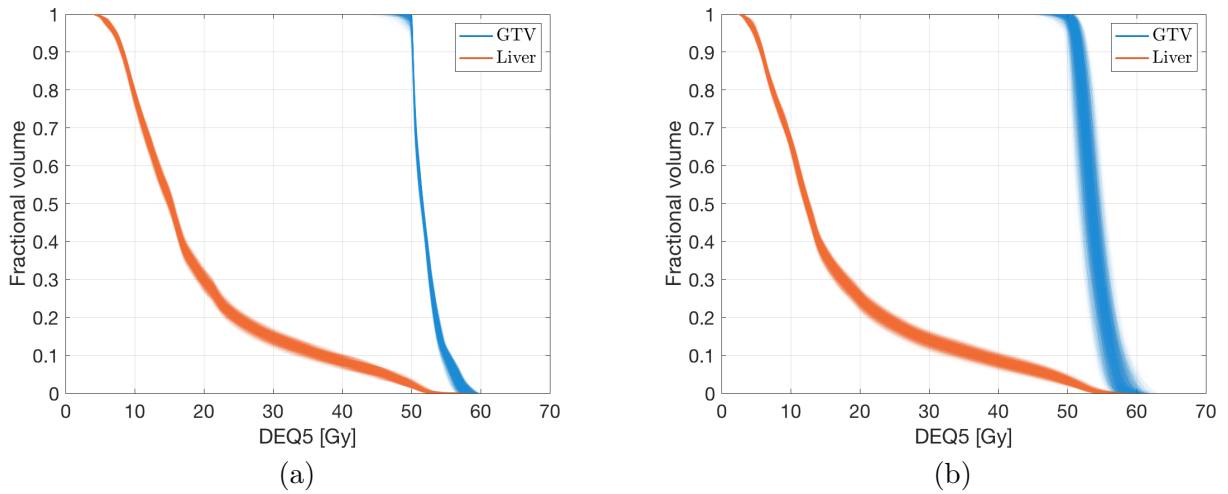


Figure B.9: A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 4 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.

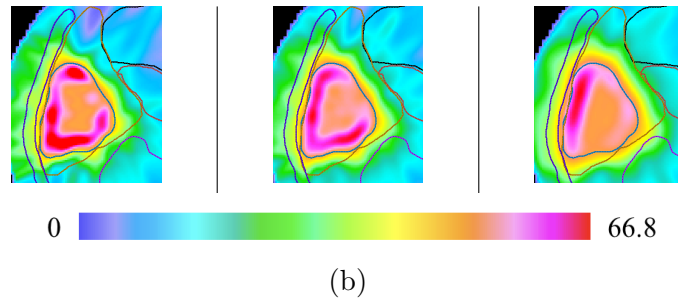
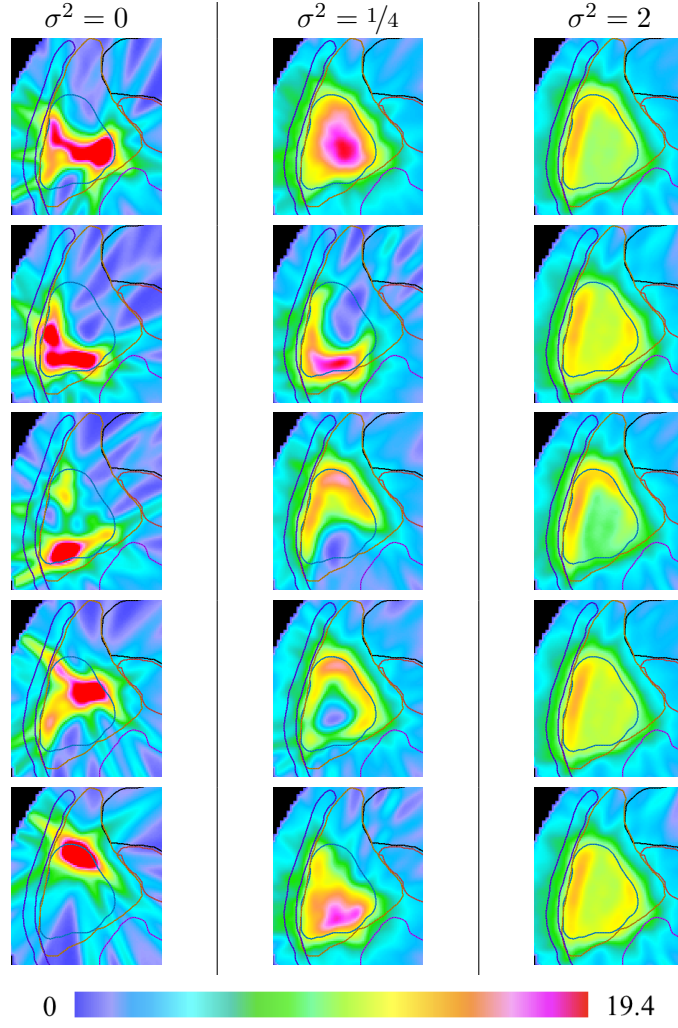


Figure B.10: (a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 5 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2 . As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.

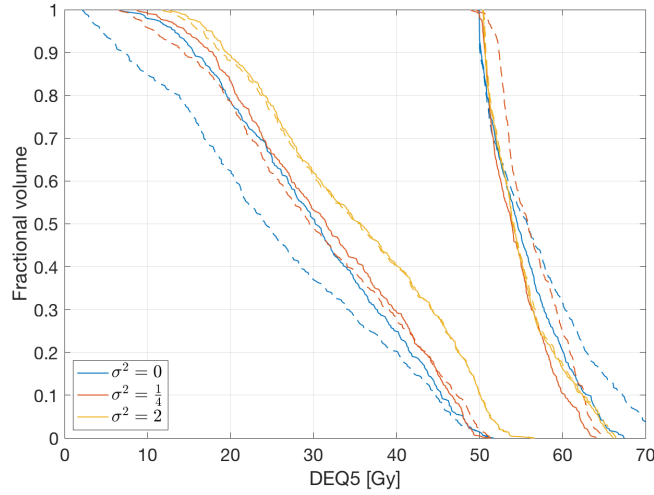


Figure B.11: Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 5, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.

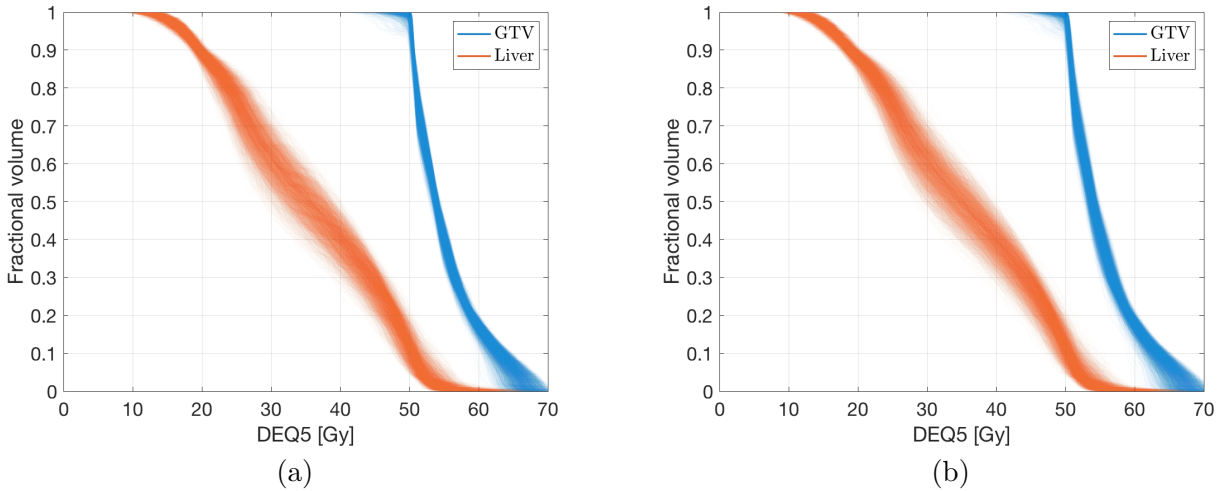


Figure B.12: A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 5 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH “cloud” is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.